

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

SYMBOL TECHNOLOGIES, INC.,
a Delaware corporation, and WIRELESS
VALLEY COMMUNICATIONS, INC.,
a Delaware corporation,

Plaintiffs/Counterclaim Defendants,

v.

ARUBA NETWORKS, INC.,
a Delaware corporation,

Defendant/Counterclaim Plaintiff.

C.A. No. 07-519-JJF

JURY DEMANDED

**DECLARATION OF ARUN CHANDRA IN SUPPORT OF
SYMBOL TECHNOLOGIES, INC.'S AND WIRELESS VALLEY COMMUNICATIONS,
INC.'S OPPOSITION TO DEFENDANT ARUBA NETWORKS, INC.'S
MOTION TO STAY PENDING RE-EXAMINATION OF THE PATENTS-IN-SUIT**

VOLUME 2 OF 2

Exhibits 12-29

OF COUNSEL

Eric J. Lobenfeld
Ira J. Schaefer
Lawrence Brocchini
Arun Chandra
HOGAN & HARTSON L.L.P.
875 Third Avenue
New York, New York 10022
(212) 918-3000

Dated April 21, 2008
861149 / 32106

Richard L. Horwitz (#2246)
David E. Moore (#3983)
POTTER ANDERSON & CORROON LLP
Hercules Plaza, 6th Floor
1313 N. Market Street
P. O. Box 951
Wilmington, Delaware 19801
Tel: (302) 984-6000
rhorwitz@potteranderson.com
dmoore@potteranderson.com

*Attorneys for Plaintiffs
Symbol Technologies, Inc. and Wireless Valley
Communications, Inc.*

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

SYMBOL TECHNOLOGIES, INC.,
a Delaware corporation, and WIRELESS
VALLEY COMMUNICATIONS, INC.,
a Delaware corporation,

Plaintiffs/Counterclaim Defendants,

v.

ARUBA NETWORKS, INC.,
a Delaware corporation,

Defendant/Counterclaim Plaintiff.

C.A. No. 07-519-JJF

JURY DEMANDED

**DECLARATION OF ARUN CHANDRA IN SUPPORT OF
SYMBOL TECHNOLOGIES, INC.'S AND WIRELESS VALLEY COMMUNICATIONS,
INC.'S OPPOSITION TO DEFENDANT ARUBA NETWORKS, INC.'S
MOTION TO STAY PENDING RE-EXAMINATION OF THE PATENTS-IN-SUIT**

I, Arun Chandra, declare as follows:

1. I am an attorney in the State of New York and an associate in the law firm of Hogan & Hartson, LLP, 875 Third Ave., New York, New York, 10022. This declaration is being submitted in support of Plaintiffs' Symbol Technologies, Inc. and Wireless Valley Communications, Inc. (hereinafter "Plaintiffs") Opposition to Defendant Aruba Networks, Inc.'s (hereinafter "Aruba") Motion to Stay Pending Re-Examination of the Patents in Suit in the above identified matter. I make this declaration as of my own personal knowledge and/or my review of the records in this action.

2. Attached here to as Exhibit 1 is a true and correct copy of the August 27, 2007 Complaint, filed in this action.

3. Attached hereto as Exhibit 2 is a true and correct copy of the October 17, 2007 Answer and Counterclaims, filed by Aruba in this action.

4. Attached hereto as Exhibit 3 is a true and correct copy of a press release describing the National Medal of Technology received by Symbol in 1999.

5. Attached hereto as Exhibit 4 is a true and correct copy of a press release describing the National Medal of Technology received by Symbol in 1999.

6. Attached hereto as Exhibit 5 is a true and correct copy of a press release describing the National Medal of Technology received by Motorola, Inc. in 2005.

7. Attached hereto as Exhibit 6 is a true and correct copy of the February 22, 2008 e-mail from Plaintiffs' counsel to Aruba's counsel, seeking agreement on a proposed scheduling order in this action.

8. Attached hereto as Exhibit 7 is a true and correct copy of the February 26, 2008 e-mail from Plaintiffs' counsel to Aruba's counsel, seeking agreement on a proposed scheduling order in this action.

9. Attached hereto as Exhibit 8 is a true and correct copy of the March 5, 2008 letter from Plaintiffs' counsel to the Court, submitting Plaintiffs' proposed scheduling order.

10. Attached hereto as Exhibit 9 is a true and correct copy of the March 7, 2008 letter from Aruba's counsel to the Court, stating Aruba's intent to seek a stay.

11. Attached hereto as Exhibit 10 is a true and correct copy of selected excerpts of the prosecution history for the '922 patent.

12. Attached hereto as Exhibit 11 is a true and correct copy of U.S. Patent No. 6,414,950.

13. Attached hereto as Exhibit 12 is a true and correct copy of U.S. Patent No. 6,665,536.

14. Attached hereto as Exhibit 13 is a true and correct copy of U.S. Patent No. 6,421,714.

15. Attached hereto as Exhibit 14 is a true and correct copy of U.S. Patent No. 6,493,679.

16. Attached hereto as Exhibit 15 is a true and correct copy of a datasheet for Aruba AP-70 Access Point.

17. Attached hereto as Exhibit 16 is a true and correct copy of a datasheet for Symbol's AP300 Access Port.

18. Attached hereto as Exhibit 17 is a true and correct copy of a datasheet for Aruba MC-2400 Mobility Controller.

19. Attached hereto as Exhibit 18 is a true and correct copy of a datasheet for Symbol's WS5100 Wireless Switch.

20. Attached hereto as Exhibit 19 is a true and correct copy of a datasheet for Aruba Mobility Management System.

21. Attached hereto as Exhibit 20 is a true and correct copy of a datasheet for Wireless Valley's EnterprisePlanner.

22. Attached hereto as Exhibit 21 is a true and correct copy of a datasheet for Wireless Valley's LANPlanner.

23. Attached hereto as Exhibit 22 is a true and correct copy of Aruba's press release, entitled *Aruba Gains Market Share and Solidifies Position as the World's Second Largest Enterprise Wireless LAN Supplier* (Sept. 7, 2007), available at <http://www.arubanetworks.com/company/news/release.php?id=33>).

24. Attached hereto as Exhibit 23 is a true and correct copy of Peter Judge, *Enterprise Wi-Fi – the Market Shifts Again*, TechWorld (Sept. 21, 2007), available at <http://techworld.com/mobility/features/index.cfm?featureid=3674>.

25. Attached hereto as Exhibit 24 is a true and correct copy of Bill Simpson, *IPO Analysis: Aruba Networks Doesn't Have Enough Under The Hood*, Seeking Alpha (April 3, 2007), available at <http://seekingalpha.com/article/31416-ipo-analysis-aruba-networks-doesn-t-have-enough-under-the-hood>.

26. Attached hereto as Exhibit 25 is a true and correct copy of the Goldman Sachs' recent research report on the Communications Technology sector.

27. Attached hereto as Exhibit 26 is a true and correct copy of the J.P. Morgan's recent research report on Aruba Networks, Inc.

28. Attached hereto as Exhibit 27 is a true and correct copy of the Jaywalk Consensus (an average of independent research providers) on Aruba Networks, Inc.

29. Attached hereto as Exhibit 28 is a true and correct copy of the PTO statistics regarding *inter partes* and *ex parte* re-examinations.

30. Attached hereto as Exhibit 29 is a true and correct copy of U.S. Patent No. 6,973,622.

I declare that the foregoing statements made by me are true. I am aware that if any of the foregoing statements made by me are willfully false, I am subject to punishment.

Dated: April 21, 2008

/s/ Arun Chandra
Arun Chandra

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

CERTIFICATE OF SERVICE

I, Richard L. Horwitz, hereby certify that on April 21, 2008, the attached document was electronically filed with the Clerk of the Court using CM/ECF which will send notification to the registered attorney(s) of record that the document has been filed and is available for viewing and downloading.

I further certify that on April 21, 2008, I have Electronically Mailed the document to the following person(s):

Frederick L. Cottrell, III
Richards, Layton & Finger
One Rodney Square
Wilmington, DE 19899
cottrell@rlf.com

Matthew D. Powers
Vernon M. Winters
Brandon C. Conard
Jason D. Kipnis
Weil, Gotshal & Manges LLP
Silicon Valley Office
201 Redwood Shores Parkway
Redwood Shores, CA 94065
matthew.powers@weil.com
vernon.winters@weil.com
brandon.conard@weil.com
jason.kipnis@weil.com

Nicholas Groombridge
Paul E. Torchia
Etai Lahav
Weil, Gotshal & Manges LLP
767 Fifth Avenue
New York, NY 10153
nicholas.groombridge@weil.com
paul.torchia@weil.com
etai.lahav@weil.com

/s/ Richard L. Horwitz
Richard L. Horwitz
David E. Moore
Potter Anderson & Corroon LLP
Hercules Plaza – Sixth Floor
1313 North Market Street
Wilmington, DE 19899-0951
(302) 984-6000
rhowitz@potteranderson.com
dmoore@potteranderson.com

EXHIBIT 12



US00665536B1

(12) **United States Patent**
Mahany

(10) **Patent No.:** **US 6,665,536 B1**
(45) **Date of Patent:** **Dec. 16, 2003**

(54) **LOCAL AREA NETWORK HAVING
MULTIPLE CHANNEL WIRELESS ACCESS**

(75) Inventor: **Ronald L. Mahany**, Cedar Rapids, IA
(US)

(73) Assignee: **Broadcom Corporation**, Irvine, CA
(US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/357,429**

(22) Filed: **Jul. 20, 1999**

Related U.S. Application Data

(63) Continuation of application No. 08/878,357, filed on Jun. 27, 1997, now Pat. No. 5,960,344, which is a continuation-in-part of application No. 08/772,895, filed on Dec. 24, 1996, now abandoned, which is a continuation-in-part of application No. 08/696,086, filed on Aug. 13, 1996, now abandoned, which is a continuation of application No. 08/238,180, filed on May 4, 1994, now Pat. No. 5,546,397, which is a continuation-in-part of application No. 08/197,392, filed on Feb. 16, 1994, now abandoned, which is a continuation-in-part of application No. 08/170,121, filed on Dec. 20, 1993, now abandoned.

(30) **Foreign Application Priority Data**

Jun. 3, 1996 (WO) PCT/US96/09474

(51) **Int. Cl.**⁷ **H04Q 7/00**

(52) **U.S. Cl.** **455/432; 455/434; 455/435**

(58) **Field of Search** 455/432, 434,
455/436, 338, 500, 444, 524, 456, 422,
466, 507, 515, 454, 554, 552; 375/131,
132, 147, 130; 370/350, 389

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,665,164 A 5/1972 Beveridge 235/460
D230,859 S 3/1974 Kurosu D26/56
3,826,900 A 7/1974 Moellering 250/568

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

EP 0194115 9/1986
FP 0262943 4/1988
FR 2389938 1/1979

(List continued on next page.)

OTHER PUBLICATIONS

"PCMCIA Primer" by John Reimer, pp. 66-67, 1995.
PCMCIA (IC) Cards, vol. 13, No. 8.
News Release "Norand Introduces Enhanced Wireless LAN Capabilities", May 16, 1995, by Norand Corporation, Cedar Rapids, Iowa, pp. 1-4.

(List continued on next page.)


Primary Examiner—Salvatore Cangialosi

(74) *Attorney, Agent, or Firm*—McAndrews, Held & Malloy, Ltd.

(57) **ABSTRACT**

A communication network having at least one access point supports wireless communication among a plurality of wireless roaming devices via a first and a second wireless channel. The access point comprises a first and a second transceiver. The first and second transceivers operate on the first and second wireless channels, respectively. Each of the plurality of wireless roaming devices are capable of communicating on the first and second wireless channel. In one embodiment, the first wireless channel is used to exchange data, while the second channel is used to manage such exchanges as well as access to the first channel. In an alternate embodiment, both channels are used to support communication flow, however the first channel supports a protocol that is more deterministic than that of the second channel. Allocation of ones of the plurality of wireless roaming devices from one channel to the next may occur per direction from the access point. It may also result from decisions made by each of the wireless roaming devices made independent of the access point. For example, a decision may be made based on the data type being transferred or based on the current channel load. Such factors may also be used by the access point for allocation determinations. In addition, allocation may be based on the type of roaming device involved, such as allocating peripherals to a slower channel.

55 Claims, 17 Drawing Sheets



US 6,665,536 B1

Page 2

U.S. PATENT DOCUMENTS

3,947,817 A	3/1976	Requa	235/472	4,842,966 A	6/1989	Omori et al.	307/150
3,976,973 A	8/1976	Martin et al.	235/472	4,845,350 A	7/1989	Shepard et al.	235/472
4,002,892 A	1/1977	Zielinski	307/150	4,850,009 A	7/1989	Zook et al.	235/375
4,005,400 A	1/1977	Engdahl	340/365 R	D303,112 S	8/1989	Desrochers	D14/100
4,017,725 A	4/1977	Roen	235/152	4,857,716 A	8/1989	Gombrich et al.	235/462
4,020,527 A	5/1977	O'Neill	16/110	4,866,646 A	9/1989	Nakamura et al.	364/709.11
4,058,838 A	11/1977	Crager et al.	358/257	4,877,949 A	10/1989	Danielson et al.	235/462
4,091,270 A	5/1978	Musch et al.	235/419	4,881,839 A	11/1989	Grimm	400/692
4,115,870 A	9/1978	Lowell	364/900	4,890,832 A	1/1990	Komaki	235/146
4,133,034 A	1/1979	Etter	364/464	4,897,532 A	1/1990	Swartz et al.	235/472
4,136,821 A	1/1979	Sugiura et al.	235/462	4,910,775 A	3/1990	Yves et al.	235/379
4,141,492 A	2/1979	Michel et al.	235/437	4,916,441 A	4/1990	Gombrich	235/380 X
4,158,130 A	6/1979	Speraw et al.	235/146	4,940,974 A	7/1990	Sojka	340/825.08
4,158,194 A	6/1979	McWaters et al.	235/454	4,942,356 A	7/1990	Ellingen et al.	361/392
4,165,554 A	8/1979	Faget	D18/7	4,953,113 A	8/1990	Chadima, Jr. et al.	364/708
4,188,103 A	2/1980	Biber et al.	354/27	4,967,188 A	10/1990	Collins	340/636
4,210,802 A	7/1980	Sakai	235/483	4,972,496 A	11/1990	Sklarew	382/13
4,247,908 A	1/1981	Lockhart, Jr. et al.	364/900	4,983,818 A	1/1991	Knowles	235/467
4,277,837 A	7/1981	Stuckert	364/900	4,984,247 A	1/1991	Kaufmann et al.	
4,282,425 A	8/1981	Chadima et al.	235/462	5,002,184 A	3/1991	Lloyd	206/305
4,322,612 A	3/1982	Lange	235/419	5,008,879 A	4/1991	Fischer et al.	370/85.2
4,385,285 A	5/1983	Horst et al.	382/3	5,012,407 A	4/1991	Finn	364/200
4,414,661 A	11/1983	Karlstrom	370/95	5,023,824 A	6/1991	Chadima, Jr. et al.	235/462
4,415,065 A	11/1983	Sandstedt	186/39	5,029,181 A	7/1991	Endo et al.	
4,418,277 A	11/1983	Tremmel et al.	235/472	5,046,066 A	9/1991	Messenger	
4,422,745 A	12/1983	Hopson	354/105	5,046,130 A *	9/1991	Hall et al.	455/456
4,423,319 A	12/1983	Jacobsen	235/419	5,049,862 A	9/1991	Dao et al.	340/706
4,460,120 A	7/1984	Shepard et al.	235/472	5,049,863 A	9/1991	Oka	340/710
4,488,679 A	12/1984	Bockholt et al.	235/469	5,057,676 A	10/1991	Komaki	235/375
4,500,776 A	2/1985	Laser	235/162	5,059,778 A	10/1991	Zouzoulas et al.	235/472
4,506,344 A	3/1985	Hubbard	364/900	5,070,536 A	12/1991	Mahany et al.	455/67
4,511,970 A	4/1985	Okano et al.	364/401	5,097,484 A	3/1992	Akaiwa	
4,519,068 A	5/1985	Krebs et al.		5,101,406 A	3/1992	Messenger et al.	
4,523,087 A	6/1985	Benton	235/379	5,117,501 A	5/1992	Childress et al.	455/11.1
4,523,297 A	6/1985	Ugon et al.	235/380	5,142,534 A	8/1992	Simpson et al.	
4,569,421 A	2/1986	Sandstedt	235/383	5,164,942 A	11/1992	Kamerman et al.	
4,570,057 A	2/1986	Chadima, Jr. et al.	364/708	5,181,200 A	1/1993	Harrison	
4,578,571 A	3/1986	Williams	235/472	5,202,817 A	4/1993	Koenck et al.	361/393
4,603,262 A	7/1986	Eastman et al.	250/566	5,216,233 A	6/1993	Main et al.	235/472
4,621,189 A	11/1986	Kumar et al.	235/472	5,218,187 A	6/1993	Koenck et al.	235/375
4,627,736 A	12/1986	Komaki	368/88	5,227,614 A	7/1993	Danielson et al.	235/380
4,628,193 A	12/1986	Blum	235/375	5,268,933 A	12/1993	Averbuch	
4,634,845 A	1/1987	Hale et al.	235/350	5,291,516 A	3/1994	Dixon et al.	
4,641,292 A	2/1987	Tunnell et al.	381/42	5,297,144 A	3/1994	Gilbert et al.	
4,654,818 A	3/1987	Wetterau, Jr.	364/900	5,321,542 A	6/1994	Freitas et al.	
4,661,993 A	4/1987	Leland et al.	455/103	5,363,401 A *	11/1994	Lucas et al.	375/131
4,689,761 A	8/1987	Yurchenco	364/708	5,363,402 A	11/1994	Harmon et al.	
4,703,161 A	10/1987	McLean	235/1	5,390,166 A	2/1995	Rohani et al.	
4,706,096 A	11/1987	Sato	235/488	5,390,206 A *	2/1995	Rein et al.	375/130
4,718,103 A	1/1988	Shojima et al.	382/13	5,404,375 A	4/1995	Kroeger et al.	
4,718,110 A	1/1988	Schaefer	455/90	5,410,141 A	4/1995	Koenck et al.	235/472
4,727,245 A	2/1988	Dobbins et al.	235/472	5,410,740 A	4/1995	Hagstrom	
4,734,566 A	3/1988	Senda et al.	235/495	5,410,752 A	4/1995	Scholefield	
D295,411 S	4/1988	Cho et al.	D14/100	5,465,392 A	11/1995	Baptist et al.	
4,743,773 A	5/1988	Katana et al.	250/566	5,469,468 A	11/1995	Schilling	375/130
4,749,353 A	6/1988	Breedlove	381/43	5,513,184 A *	4/1996	Vannucci	370/350
4,752,965 A	6/1988	Dunkley	235/379	5,546,397 A	8/1996	Mahany	
4,758,717 A	7/1988	Shepard et al.	235/472	5,551,060 A *	8/1996	Fujii et al.	455/444
4,760,248 A	7/1988	Swartz et al.	235/472	5,561,845 A *	10/1996	Bendixen et al.	455/434
4,761,822 A	8/1988	Maile	455/82	5,644,576 A *	7/1997	Bauchot et al.	375/131
4,766,300 A	8/1988	Chadima, Jr. et al.	235/472	5,655,219 A *	8/1997	Jusa et al.	455/338
4,768,230 A	8/1988	Viebrantz et al.	455/603	5,734,645 A *	3/1998	Raith et al.	455/466
4,773,032 A	9/1988	Uehara	364/900	5,748,621 A	5/1998	Masuda et al.	370/337
4,793,812 A	12/1988	Sussman et al.	235/472	5,768,267 A *	6/1998	Raith et al.	370/329
D299,234 S	1/1989	Kajita	D14/100	5,790,587 A *	8/1998	Smith et al.	375/147
4,825,057 A	4/1989	Swartz	235/472	5,844,893 A	12/1998	Gollnick et al.	370/329
4,831,275 A	5/1989	Drucker	235/472	5,960,344 A *	9/1999	Mahany	455/432
4,835,372 A	5/1989	Gombrich	235/375	6,118,824 A	9/2000	Smith et al.	375/259
4,836,256 A	6/1989	Meliconi	206/523	6,138,010 A	10/2000	Rabe et al.	455/426
4,837,858 A	6/1989	Ablay et al.	455/34	6,185,205 B1 *	2/2001	Sharrit et al.	370/389

US 6,665,536 B1

Page 3

FOREIGN PATENT DOCUMENTS

GB	1572962	8/1980
GB	2201125	8/1988
JP	58176792	10/1983
JP	58-211261	12/1983
JP	2-144681	6/1990
WO	8700659	1/1987
WO	8707106	11/1987

OTHER PUBLICATIONS

News Release "Ositech Communications Inc. Will Release Next Month the Jack of Diamonds, A Type II PC Card", dated May 31, 1995.

Roger C. Alford, "The PCMCIA Redefines Portability", Under The Hood, 1995.

Gary Legg, "Special Components Simplified Interface to PCMCIA cards", Jun. 10, 1993 EN., pp. 61, 63-64, 66 and 68.

Carole Patton, "Radio Transceiver Circuit Card Acts as Wireless Modem in PCs", Info World, Mar. 28, 1988, p. 18.

John R. Gyorki, "New Deal For Smart Cards", Machine Design, Jun. 1, 1993, pp. 38-40, 42 and 44.

* cited by examiner

U.S. Patent

Dec. 16, 2003

Sheet 1 of 17

US 6,665,536 B1




Fig. 1

U.S. Patent

Dec. 16, 2003

Sheet 2 of 17

US 6,665,536 B1




Fig. 2




Fig.3

U.S. Patent

Dec. 16, 2003

Sheet 4 of 17

US 6,665,536 B1




Fig. 4




Fig.5

U.S. Patent

Dec. 16, 2003

Sheet 6 of 17

US 6,665,536 B1




Fig. 6

U.S. Patent

Dec. 16, 2003

Sheet 7 of 17

US 6,665,536 B1




Fig. 7a

U.S. Patent

Dec. 16, 2003

Sheet 8 of 17

US 6,665,536 B1




Fig. 7b

U.S. Patent

Dec. 16, 2003

Sheet 9 of 17

US 6,665,536 B1




Fig. 8

U.S. Patent

Dec. 16, 2003

Sheet 10 of 17

US 6,665,536 B1




Fig.9

U.S. Patent

Dec. 16, 2003

Sheet 11 of 17

US 6,665,536 B1




Fig.10




Fig. 11

U.S. Patent

Dec. 16, 2003

Sheet 13 of 17

US 6,665,536 B1




Fig.12

U.S. Patent

Dec. 16, 2003

Sheet 14 of 17

US 6,665,536 B1




Fig. 13

U.S. Patent

Dec. 16, 2003

Sheet 15 of 17

US 6,665,536 B1




Fig. 14a

U.S. Patent

Dec. 16, 2003

Sheet 16 of 17

US 6,665,536 B1




Fig. 14b

U.S. Patent

Dec. 16, 2003

Sheet 17 of 17

US 6,665,536 B1




Fig. 15

US 6,665,536 B1

1

**LOCAL AREA NETWORK HAVING
MULTIPLE CHANNEL WIRELESS ACCESS****CROSS-REFERENCE TO RELATED
APPLICATIONS**

The present application is a continuation of U.S. application Ser. No. 08/878,357 filed Jun. 27, 1997 now U.S. Pat. No. 5,960,344, which is a continuation-in-part of U.S. application Ser. No. 08/772,895 filed Dec. 24, 1996, abandoned, which is a continuation-in-part of U.S. application Ser. No. 08/696,086 filed Aug. 13, 1996, abandoned, which is a continuation of U.S. application Ser. No. 08/238,180 filed May 4, 1994, now issued as U.S. Pat. No. 5,546,397, which is a continuation-in-part of U.S. application Ser. No. 08/197,392 filed Feb. 16, 1994, abandoned, which is a continuation-in-part of U.S. application Ser. No. 08/170,121 filed Dec. 20, 1993, abandoned.

The U.S. application Ser. No. 08/772,895 filed Dec. 24, 1996, also claims priority to PCT application Ser. No. PCT/US96/09474, filed on Jun. 3, 1996.

All of the aforementioned applications are hereby incorporated herein by reference in their entirety. In addition, U.S. Pat. No. 5,425,051 issued Jun. 13, 1995 to Ronald L. Mahany is also hereby incorporated herein by reference in its entirety.

BACKGROUND**1. Technical Field**

The present invention relates generally to access points used in wireless local area networks, and more specifically to an access point which includes multiple wireless adapters.

2. Related Art

Wireless local area networks (WLAN's) use radio frequency transmissions to communicate between roaming computer devices and access points (or base stations). The access points are connected to an infrastructure that electronically connects all of the access points to a host system. The wired infrastructure and the access points make up an information distribution network used for the transfer of information and for communications.

In a wireless networking environment, various types of devices may need to communicate within a given area. When incompatibilities between device types arise, the wireless infrastructure must accommodate the various device types. Accommodating the different device types in a single infrastructure is generally difficult to accomplish. Further, devices within the wireless networking environment typically communicate differing types of data, each with its own priority and bandwidth requirements. Accommodating the various types of data with their related priorities often could not be accomplished by prior devices due to bandwidth limitations, conflicting priorities and incompatible standards within the wireless network.

In prior WLANs, a first wireless terminal that desired to communicate with a base station often could not detect transmissions from a second wireless terminal currently engaged in ongoing communication with the access point. As a result, the wireless terminal often initiated transmissions that collided with the ongoing communications. Operation of this type is referred to as a "hidden terminal" situation. To solve the hidden terminal situation, some prior base stations were configured with a second transmitter for delivering a carrier signal on a "busy channel" whenever the base station was engaged in communication on the "data channel." All terminals were also fitted with a second

2

receiver, tuned to the busy channel, and required to check the busy channel before initiating communication on the data channel. However, the additional power required, bandwidth used, hardware needed and associated cost made the busy channel solution undesirable for most applications.

Some prior WLANs attempted to solve operational difficulties by simply increasing the transmission capacity available on the infrastructure. Such expansion temporarily decreased conflicts in operation of the WLANs. However, the infrastructure, which is expensive to install, typically became overloaded quickly resulting in the same or similar problems.

SUMMARY OF THE INVENTION

The present invention is directed to communication network that supports communication within a premises. The communication network comprises an access point, a plurality of wireless roaming devices, a first wireless communication channel, and a second wireless communication channel. The access point itself comprises a first processing circuit, a first radio transceiver coupled to the first processing circuit, and a second radio transceiver coupled to the first processing circuit. Each of the plurality of wireless roaming devices comprising a second processing circuit, a third radio transceiver and a radio receiver. Therein, the first wireless communication channel that supports communication flow via the communication network, while the second wireless communication channel is used to manage the flow of communication through the first wireless communication channel. In addition, the first and third radio transceivers are operable on the first wireless communication channel, while the second radio transceiver and the radio receiver are operable on the second wireless communication channel.

The communication network also supports various other aspects of the present invention. For example, the access point may further comprise a wired communication interface circuit coupled to the first processing circuit. Selective participation on the first and second communication channels may also provide further benefits. In one embodiment, each of the plurality of wireless roaming devices utilizes the radio receiver on the second wireless communication channel before participating with the third radio transceiver on the first wireless communication channel. In another, each utilizes the radio receiver on the second wireless communication channel to gain access with the third radio transceiver on the first wireless communication channel. Each may also or alternatively utilize the second wireless communication channel to identify ongoing communication on the first wireless communication channel to, perhaps, provide an indication as to when channel capacity may become available.

Other aspects may be found in an alternate communication network which also supports communication within a premises. This communication network comprises an access point, first and second wireless communication channels and plurality of wireless roaming devices. The first wireless communication channel has first communication flow characteristics, while the second wireless communication channel has second communication flow characteristics. The first and second radio transceivers participate on the first and second wireless communication channels, respectively. Therein, each of the plurality of wireless roaming devices comprises a second processing circuit and means for selectively participating on the first and second wireless communication channels.

The access point may also comprise a wired communication interface circuit coupled to the first processing circuit

US 6,665,536 B1

3

that may itself comprise a first and a second microprocessor. Additionally, at least one of the plurality of wireless roaming devices may participate on the first wireless communication channel while the other of the plurality of wireless roaming devices participate on the second wireless communication channel. Although the at least one of the plurality of wireless roaming devices may participate on the first wireless communication channel as directed by the access device, other variations and combinations are also possible. For example, at least one of the plurality of wireless roaming devices may participate on the first wireless communication channel to exchange a specific type of data, and/or may participate based on current channel conditions. Such participation may be based the fact that, in some embodiments, the second wireless communication channel is more deterministic than the first wireless communication channel.

In any of the aforementioned embodiment, the communication network may comprise at least a second access point. Other variations and aspects of the present invention will become apparent to ones of ordinary skill in the art after reviewing the entire specification and drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic representation of a high reliability access point in accordance with the present invention.

FIG. 2 is a schematic representation of another high reliability access point of the present invention utilizing an antenna diversity scheme at each wireless adapter.

FIG. 3 is a representation of a distribution network for a wireless LAN system utilizing high reliability access points.

FIG. 4 is a schematic representation of a high reliability access point with a backup power supply.

FIG. 5 is a schematic representation of a remote high reliability access point connecting to the distribution network.

FIG. 6 is block diagram illustrating an embodiment of an access point built in accordance with the present invention which includes two radios and a wired network interface, a first one of the radios operable on a first channel and a second one of the radios operable on a second channel.

FIG. 7a is a block diagram illustrating an embodiment of a portable data terminal according to the present invention, the portable data terminal having a single PCMCIA card that contains two radios, a first one of the radios operable on the first channel and a second one of the radios operable on the second channel.

FIG. 7b is a block diagram illustrating an alternative embodiment of the portable data terminal of FIG. 7a, wherein the single PCMCIA card includes a single radio operable on the first channel and the second channel and controlled by the processing circuitry.

FIG. 8 is a block diagram illustrating an alternative embodiment of a portable data terminal according to the present invention, the portable data terminal having a single PCMCIA card that contains a multi-channel wireless transceiver and a wired network interface.

FIG. 9 is a diagram illustrating a communication system built and operating according to the present invention, the communication system including at least one access point having multiple radios, portable terminals having multiple radios and portable terminals having multi-channel radios.

FIG. 10 is a diagram illustrating a communication system built and operating according to the present invention wherein one of the access points facilitates communication between portable terminal units operating on different channels within its cell by routing communication between two of its radios.

4

FIG. 11 is a block diagram illustrating an embodiment of a communication system according to the present invention wherein an access point uses a dedicated control/busy channel transmitter to manage transmissions between the access point and a plurality of roaming portable data terminals within its cell.

FIG. 12 is a drawing illustrating advantageous operation of the access device and portable data terminals of FIG. 11 when two roaming terminals encounter hidden terminal conditions.

FIG. 13 is a block diagram illustrating an alternate embodiment of the communication system of the present invention wherein an access point includes a dedicated control/busy channel transceiver and roaming data terminals communicate with the access point using either frequency nimble multi-channel transceivers or dedicated control/busy channel transceivers.

FIG. 14a is a block diagram illustrating a communication system of the present invention wherein access points and portable data terminals operate on a deterministic first channel and a non-deterministic second channel and the system routes communications on the channels based upon system conditions.

FIG. 14b is a diagram illustrating operation of a communication system of the present invention having both wired and wireless communication capability that includes at least one access point providing communication over a deterministic, time bounded first channel and a non-deterministic, contention access second channel.

FIG. 15 is a diagram illustrating the use of the access points and portable data terminals of FIG. 14a wherein the system routes various transmissions within the network system according to system conditions such as channel activity, data type and data priority.

DETAILED DESCRIPTION

Referring now to the drawings wherein like reference numerals designate identical or corresponding parts throughout the several views, FIG. 1 shows a high reliability access point 10 built in accordance with the present invention. An access point is a base station on a wireless local area network with which roaming portable or mobile computer devices can connect and communicate. The access point is typically part of an overall distribution network which is connected to a host computer or entire computer local area network (LAN). The access points and the infrastructure make up the distribution network and allow for communications between the roaming computer devices and the host computer or entire computer local area network (LAN).

A high reliability access point 10 of the present invention includes a central processing unit CPU processor 13 and at least two wireless adapters 15 and 16. Each of the wireless adapters 15 and 16 include a radio 17 and 18, a media access control (MAC) processor 19 and 20 and an antenna 21 and 22, respectively. The radios and antennas are used for RF transmission and reception. The MAC processor controls low level protocol functions including controlling the operation of the radio, radio channel, error control, e.g., ARQ or Selective Response, and communication with the CPU processor 13. The CPU processor 13 controls the high-level communications protocol functions and controls the interface 25 between the high reliability access point 10 and the infrastructure 26. In a preferred embodiment there is a PCMCIA standard interface between the wireless adapters and the access point.

The distribution network is comprised of all of the access points and the infrastructure which connects all of the access

US 6,665,536 B1

5

points. A host computer or an entire host network is connected to the distribution network. The distribution network allows computer devices to communicate with the host computer or host network.

The division between what is high level protocol, and thus handled by the CPU processor, and what is low level protocol, and thus handled by the MAC processor, can vary greatly depending upon the intelligence level of the MAC processor. In a preferred embodiment, the infrastructure conforms to an industry standard wired LAN such as Ethernet. The MAC processor can be made very intelligent and therefore capable of handling a great deal of radio specific protocol. On the other hand, the MAC processor can be minimally intelligent and handle only the most basic protocol functions allowing the CPU processor to handle the majority of the protocol functions.

Utilizing multiple wireless adapters in a single access point, as well as incorporating independent intelligence and low level protocol responsibility into each wireless adapter, yields several significant advantages. The examples depicted in FIGS. 1-4 show access points using only two wireless adapters per access point. Utilizing two wireless adapters in the manner discussed below will greatly increase the reliability of a particular access point, as well as increase the reliability of the entire distribution network. Access points could use more than two wireless adapters and the utilization of the multiple wireless adapters would be similar to the implementation describes using only two wireless adapters with addition protocol being required to handle the increased redundancy and to allow for more sophisticated self monitoring.

Referring still to FIG. 1, the CPU processor 13 can designate the RF address to which each wireless adapter 15 and 16 is to respond. The CPU processor 13 can, but need not, assign the same address to each wireless adapter. Therefore, in one configuration, the CPU processor 13 can designate that each of the wireless adapters 15 and 16 respond to the address assigned to that access point 10. Designated as such, both radios 17 and 18 will be operating simultaneously on the same channel. In a frequency hopping system, both radios 17 and 18 would be operating on the same hopping sequence, and be mutually synchronized to that hopping sequence.

Accordingly, the wireless adapters 15 and 16 are configured to receive incoming transmission from roaming computer devices within range. As both wireless adapters 15 and 16 receive the transmission, each adapter can evaluate the quality information to the CPU processor 13. The CPU processor 13 uses the quality information to determine which wireless adapter is receiving the higher quality signal. The CPU processor 13 will then typically choose to receive the incoming transmission on the wireless adapter with the higher signal quality and respond using the same adapter.

The antennas 21 and 22 can be positioned to allow the access point 10 to implement an antenna diversity scheme which will help reduce the negative effects caused by multipath interference. Antenna diversity can be accomplished in various ways. For example, the antennas can be placed sufficiently far apart, typically greater than a quarter wavelength apart, or the antennas can be positioned at a 90 degree angle with respect to each other to create a polarization antenna diversity scheme.

With an antenna diversity scheme in place, the signal from a wireless computer device will be received differently on each antenna due to multipath signal propagation. Therefore, each wireless adapter may receive a signal of a different

6

quality. The CPU processor 13 can choose which wireless adapter to use based upon the quality of the received signal. Each wireless adapter includes the capability of measuring signal quality and only good messages will be forwarded on to the CPU processor 13. The quality can be appended to the message or can be presented to the CPU in a memory register.

Referring now to FIG. 2, another high reliability access point 20 built in accordance with the present invention is shown. In this embodiment, in addition to having an antenna diversity scheme at the access point level, there is an antenna diversity scheme at the wireless adapter level. Each wireless adapter 15 and 16 includes at least two antennas 21 and 23, 22 and 24, respectively positioned to create an antenna diversity scheme. Thus for the wireless adapter 15 the antennas 21 and 23 are either positioned sufficiently far apart, more than a quarter wavelength, or the antennas 21 and 23 are positioned in an asymmetrical or orthogonal manner to provide polarization diversity. The antennas 22 and 24 for the wireless adapter 16 are placed in a similar manner.

In this embodiment, an incoming signal is received on both antennas 21 and 23 of the wireless adapter 15. The MAC processor 19 then determines the quality if the signal coming in on each of the antennas 21 and 23 connected to the wireless adapter 15. Based upon the signal quality information, the MAC processor 19 will choose which of the antennas 21 and 23 to use to receive the incoming transmission. The MAC processor will also forward the signal quality information regarding the selected antenna to the CPU processor 13. The wireless adapter 16 will perform a similar process and forward the signal quality information for its best antenna to the CPU processor 13. The CPU processor 13 can then determine which wireless adapter is receiving the highest quality signal and use that wireless adapter to receive the incoming transmission and respond to the transmitting station.

When a high reliability access point wishes to transmit a message, such as an acknowledgment of a received message, to a roaming computer device, the CPU processor 13 will utilize the received quality signal information to determine which wireless adapter to use to send the message. Likewise, if the wireless adapter is utilizing an antenna diversity scheme it will also select the most appropriate antenna for transmitting a message.

While one of the wireless adapters is transmitting, the other wireless adapter can operate as a promiscuous listener to determine if the correct message is being sent. For example, referring to FIG. 1, if the CPU processor 13 is sending a message to a roaming computer device via wireless adapter 16, wireless adapter 17 can operate in the receive mode and monitor the message being sent by the wireless adapter 16. This provides a local loop back capability and allows the access point to perform self-monitoring. If the CPU processor 13 determines that one of the wireless adapters is not operating correctly, the malfunctioning wireless adapter can be disabled. Additionally, the CPU processor 13 can then send a message to the system management portion of the host network via the infrastructure 26 that it has a defective wireless adapter and repairs are needed.

Referring again to the configuration in which each of the wireless adapters is listening on the same channel, another advantage achieved by this configuration is the ability to receive two concurrent messages. In an access point that only contains one adapter, this situation will cause a colli-

US 6,665,536 B1

7

sion and neither message will be received. In a high reliability access point built in accordance with the present invention, it is possible that the one wireless adapter will be able to receive one of the messages while the other wireless adapter receives the other, due to multipath fading at each of the wireless adapter antennas.

Referring now to FIG. 3, a portion of a distribution network 30 utilizing high reliability access points is shown. The distribution network 30 includes an infrastructure 33 and two high reliability access points 35 and 36. Access point 35 includes a CPU processor 37 and two wireless adapters 38 and 39. Access point 36 includes a CPU processor 41 and two wireless adapters 42 and 43. In the present example, a break 45 in the infrastructure 33 has occurred. Access point 35 is upstream to the break with respect to the host computer network and thus is not immediately affected by the break 45. However, access point 36 is downstream to the break 45 and therefore is no longer connected to the host computer network.

When a situation like this occurs, the downstream access point 36 will begin attempting to communicate with an upstream access point using wireless communication. In this example, the upstream access point is access point 35. However, the communication need not be with the access point immediately upstream, the only requirement is that it be with an access point which is upstream with respect to the break. The host computer network or other access points will previously have shared the logistic and address information concerning all of the access points to each access point in the distribution network.

Once communications with an upstream access point 35 is established, each access point 35 and 36 will dedicate one of its wireless adapters 39 and 42, respectively to provide a wireless repair of the break 45 in the infrastructure 33. When this happens, the CPU processor for each of the access points will instruct the dedicated wireless adapter to change so that it is no longer operating on the same channel as the other adapter in the access point. A communication channel between access points is established. The dedicated wireless adapters 39 and 42 will no longer be used to transmit or receive information to or from roaming computer devices. However, the non-dedicated wireless adapters 38 and 43 will communicate with the roaming computer devices. Once the top priority of re-establishing communications between all of the access points in the distribution network 30 and the host computer network has been accomplished, the access points can then send a message to the system management portion of the host computer network detailing where the break (or breaks) exists.

It is conceivable that the distribution network could lose its entire infrastructure. In this case, each of the high reliability access points would dedicate one of its wireless adapters to network infrastructure communications while retaining one of its wireless adapters for communication with roaming computer devices. Using the same technique described above, a temporary or remote access point could be established that, intentionally, is not connected to the infrastructure. This configuration is discussed below in greater detail with reference to FIG. 5. The use of directional gain antennas for the dedicated wireless adapter would allow the temporary or remote access point to be positioned a considerable distance from the infrastructure.

Referring now to FIG. 4, a high reliability access point 50 with a backup power supply 52 is shown. Typically, the access point will be wired to an external power source 54 such as a wall outlet. However, there is a great desire that if

8

power is lost that the distribution network not shut down since the roaming computer devices will normally not be dependent upon the external power source 54. In this embodiment of the present invention, the back-up power source 52 is wired in parallel with respect to the external power source 54. Thus, if the external power source 54 fails, the access point 50 will not lose power.

Referring now to FIG. 5, a remote access point 70 is shown connecting to the infrastructure 33 by means of dedicated wireless adapters 62 and 72. The access point 70 is not hard wired to the infrastructure 33. Therefore, the access point dedicates one of its wireless adapters 72 to network infrastructure communication. The other wireless adapter 74 continues to communicate with roaming computer devices within the range of the access point 70. An access point 60 that is hard wired into the infrastructure 33 dedicates one of its wireless adapters 62 to network infrastructure communication and establishes a link between the infrastructure 33 and the remote access point 70. The access point 60 can continue to service the roaming computer devices within its range through the wireless adapter 64.

The hard wired access point 60 that is used to connect the remote access point 70 to the infrastructure need not be the access point that is physically closest to the remote access point 70. Use of the directional antenna would allow a remote access point to establish communication with virtually any of the access points that are hard wired to the infrastructure. Additionally, several remote access points could establish wireless infrastructure communication by each dedicating one of their wireless adapters. In this arrangement, only one of the remote access points need be in communication with a wired access point. All other remote access points could establish communication with the host computer network via the remote access point in communication with a wired access point.

FIG. 6 is block diagram illustrating an embodiment of an access point 600 built in accordance with the present invention capable of communicating with wireless devices in its cell on both a first channel and a second channel. The access point 600 thus includes a first radio 616 operating on a first channel and a second radio 608 operating on a second channel. The access point also includes a processing unit 612 and additional circuitry 614, both of which couple to the first radio 616, the second radio 608 and a wired Ethernet transceiver through a bus interface 610. The wired transceiver 606 allows the access point 600 to access a wired LAN backbone 622 to which various other system components may connect. The wired LAN backbone may include, for example, an ethernet network, a token-ring network or an asynchronous transfer mode (ATM) network among other network types. In any such case, the wired transceiver 606 facilitates communication between the access point 600 and devices coupled to the wired LAN backbone 622.

The blocks illustrated in FIG. 6 are simplified for exemplary purposes, and it will be understood by one skilled in the art that an access point 600 according to the present invention is not limited to the block circuitry shown in FIG. 6. In another embodiment, the access point 600 may contain additional transceivers for communicating on other channels, over other mediums and in other networks as well.

The first channel radio 616 couples to first antenna 618 while the second channel radio 608 couples to second antenna 620. The antennas 618 and 620 may either be protruding or non-protruding antennas, depending upon system requirements. The first channel radio 616 and the second channel radio 608 operate independently to form a

US 6,665,536 B1

9

first communication cell and a second communication cell, respectively. When a radius of the first communication cell substantially equals a radius of the second communication cell, the cells substantially overlay one another. However, when the radii of the communication cells differ, the larger cell fully overlays and extends beyond the smaller cell. The first and second channels may operate using different frequencies, modulation schemes and code spreading schemes. The selections of such operational channel variations depend on overall system constraints, yet should result in two independent channels that do not interfere with one another unacceptably.

The bus interface 610 isolates the processing unit 612 and the additional circuitry 614 of the access point 600 from the operating characteristics of the radios 616 and 608 and the wired transceiver 606. Thus, communication with any of the transceivers can be accommodated by general circuitry and software routines of the access point 600. In one embodiment, the bus interface 610 is a PCI bus interface with the first channel radio 616, second channel radio 606 and wired transceiver compatible with PCI bus standards. However, in other embodiments, differing interface standards may be employed.

In operation, the processing unit 612 is programmed with the network configuration to route communications through the first channel radio 616, the second channel radio 608 and the wired transceiver 606. However, roaming portable units may alter the network configuration as they move between cells. Thus, the access point 600 periodically polls devices within its communication cell to update the network configuration. Updates are entered and forwarded for other units in the system.

Incoming messages received via the wired transceiver 606 may be stored, displayed and routed via the first channel radio 616 or routed via the second channel radio 608 to portable data terminals or other wireless devices operating within the cell(s) of one or more of the access point 600. Similarly, an incoming message on one of the radios 616 or 608 may be stored, displayed, routed through one of the radios 616 or 608 or routed through the wired transceiver 606, depending upon the message destination and type.

By providing routing within the access point 600 between the first channel radio 616 and the second channel radio 608, message delivery is expedited. Further, as will be described herein, by providing two radios in various access points, fewer cells may be required to adequately service a premises such as a factory. Moreover, when one of the radios is employed to provide control within a cell while the other radio provides primary communication within the cell, collisions between devices may be eliminated. Still further, when one of the radios provides a deterministic communication path while another one of the radios provides a non-deterministic communication path, data and message transmissions within the network may be controlled to satisfy bandwidth requirements of the various devices within the system. It may be preferable to utilize a deterministic communication path for some types of communications such as telephony video or real-time data transfer, for example. However, when the preferred deterministic path is unavailable for some reason, the alternative non-deterministic path may still be used.

The access point 600 may synchronize transmissions on the first channel radio 616 and the second channel radio 608 to avoid unacceptable conflicts between transmissions on one radio and receipts on the other radio. In this fashion, unacceptable conflicts are minimized.

10

FIG. 7a is a block diagram illustrating an embodiment of a portable data terminal 720 according to the present invention, the portable data terminal having a single PCMCIA card that contains two radios. In particular, the portable data terminal 720 contains terminal circuitry 722 that includes processing circuitry 726, conventional terminal circuitry 728 and interface circuitry 730. The interface circuitry 730 provides a PCMCIA interface for receiving PCMCIA cards of various functionality. Terminal circuitry 722 is well known and can be found in conventional portable or hand held computing devices.

Via the interface circuitry 730, the portable data terminal 720 accepts PCMCIA cards. As illustrated, the PCMCIA card inserted constitutes a communication module 724 that provides wireless access on two channels. Specifically, the communication module 724 comprises processing circuitry 732, first channel radio 735, second channel radio 734 and interface circuitry 744. The first channel radio 735 communicates via first antenna 737 while the second channel radio 734 communicates via second antenna 738. Configured and operable in this manner, the portable data terminal 720 may communicate with the access point 600 of FIG. 6 on either the first channel or the second channel.

Independent of whether the first channel radio 735 or the second channel radio 734 is used, the processing circuitry 726 delivers and receives data and messages via the interface circuitry 730 in the same manner and format, i.e., the interface circuitry 730 supports a common communication interface and protocol. The processing circuitry 732 of the communication module 724 receives data and messages via the interface circuitry 744. The processing circuitry 732, including a DSP 742, participates to assist in wireless communication via both the first channel radio 735 and the second channel radio 734. Thus, the module 724 not only saves on PCMCIA slots, but also saves costs and increases reliability by sharing common circuitry resources. In particular, the first channel radio 735 and second channel radio 734 share the interface circuitry 744 and processing circuitry 732 which includes the DSP 742. In another embodiment of the portable data terminal 720, a PCMCIA compatible wired network adapter could be installed which would also share some of the common circuitry resources.

FIG. 7b is a block diagram illustrating an alternative embodiment of a portable data terminal 748 that receives a single PCMCIA card having a radio 750 that includes two separate radio units. As contrasted to the dual radio design of the portable data terminal 720 of FIG. 7a, the radio 750 of the portable data terminal 748 of FIG. 7b operates on both the first channel and the second channel. The radio 750 is coupled to antenna 756 and controlled by processing circuitry 752 that includes digital signal processing circuitry 754. The radio 750 includes a first radio unit operable on the first channel and a second radio unit operable on the second channel with the radio units sharing some common components.

The processing circuitry 752 may control operation of the radio 750 in a simplex fashion such that the radio 750 operates on the first channel as required and operates on the second channel as required. Because the radio 750 may include circuitry shared by the radio units, the radio 750 may only operate on one channel at a given time. By multiplexing its operation over time, however, the radio 750 provides sufficient coverage on the channels at a reduced cost. Other components of the portable data terminal of FIG. 7b were previously described with reference to FIG. 7a and will not be further described herein.

FIG. 8 is a block diagram illustrating an alternative embodiment of a portable data terminal 800 according to the

US 6,665,536 B1

11

present invention, the portable data terminal **800** having a single PCMCIA card that contains a multi-channel (or multi-mode) wireless transceiver **739** and a wired network interface **736** (or modem transceiver). The portable terminal **800** includes terminal circuitry **722** and a module **802** including various components previously described with reference to FIG. 7a. The terminal circuitry **722** includes processing circuitry **726**, conventional terminal circuitry **728** and interface circuitry **730**. The communication module **802** includes processing circuitry **732**, the multi-mode wireless transceiver **739**, the wired modem transceiver **736** and interface circuitry **744**. When in use, the wired modem transceiver **736** interfaces via a jack **740** to a telephone line (not shown). Similarly, the wireless multi-mode transceiver **739** communicates via an antenna **741**.

The processing circuitry **732** of the communication module **802** receives data and messages via the interface circuitry **744**. If the modem transceiver **736** is being used, the processing circuitry **732** appropriately (de)segments and (de)compresses the data/messages utilizing a digital signal processor (DSP) **742**. Otherwise, the processing circuitry **732**, including the DSP **742**, participate to assist in wireless communication via the multi-mode transceiver **739**. Thus, the module **802** not only saves on PCMCIA slots (as required when a conventional radio card and a conventional modem card are both being used), but also saves costs and increases reliability by sharing common circuitry resources.

The multi-mode transceiver **739** is frequency nimble and may operate in various modes, such as those that may be used with a frequency spreading scheme such as those described in U.S. Pat. No. 5,425,051 issued Jun. 13, 1995 to Ronald L. Mahany, which is incorporated herein by reference. Thus, the multi-mode transceiver **739** may operate on both the first channel and the second channel and communicate with the access point **600** of FIG. 6 on either the first channel or the second channel. As will be further described herein, operation on differing channels may be employed to reduce installed system component requirements, to alleviate various potential interfering operating conditions and to more efficiently route data and messages within the wireless local area network.

FIG. 9 is a diagram illustrating a communication system **900** built and operating according to the present invention. The communication system includes an access point **902** operating on two channels and access points **904** and **906** operating on a single channel. Each of the access points **902**, **904** and **906** connects to a wired LAN backbone **908** to facilitate wired communication between the access points and computer systems **910** and **912** connected to the wired LAN backbone **908**.

Access point **902** includes both a first channel radio and a second channel radio. In the embodiment illustrated, the first channel radio creates a first channel cell **930** extending with a first channel radius about the access point **902**. The second channel radio of the access point **900** creates a second channel cell **932** extending with a second channel radius about the access point **902**. As illustrated, the second channel cell **932** has a larger radius than the radius of the first channel cell **930**. To create the relatively larger cell, the second radio may operate at a higher power, operate at a lower data rate or operate in another differing manner to create the relatively larger cell.

Access points **904** and **906** generate first channel cells **934** and **936**, respectively. Portable data terminals **920**, **922**, **924** and **926** and scanning unit **918** communicate with the various access points **902**, **904** and **906** and roam about the

12

communication system **900**, potentially moving from cell to cell. Other devices, such as stationary printers **914** and **916** typically remain within one cell of the communication system **900**. In the embodiment illustrated, portable data terminals **920** and **926** include multi-mode radios while portable data terminals **922** and **924** include both a first channel radio and a second channel radio. However, in other embodiments, some of the portable data terminals may only on one of the channels.

As illustrated, terminal **922** includes a first channel antenna **940** and a second channel antenna **942** while access point **902** includes both a first channel antenna **944** and a second channel antenna **946**. Thus, whenever the terminal **922** roams within the second channel cell **932**, the terminal **922** communicates via a second channel radio and second channel antenna **942**. Further, terminal **950** may communicate with access point **902** on the first channel via its first channel antenna **954** when resident within the first channel cell **903**. Finally, terminal **952**, having a single radio operable the first channel via antenna **956** may communicate with access point **902** on the first channel when resident within the first channel cell **932**.

As shown, portable data terminal **922** resides both within the first channel cell **930** and the channel cell **932** generated by access point **902**. Thus, the portable data terminal **922** may communicate with access point **902** on either the first channel or the second channel. However, printer **916** and scanning unit **918** reside only within the second channel cell **932** generated by access point **902** and must communicate with the access point **902** on the second channel.

Print data originating at computer **910** and intended for printer **916** travels from computer **910**, through the wired LAN backbone **908** to access point **902** and across first channel cell **932** to the printer **916**. During this transmission, the data is routed through the wired LAN backbone and the wireless network based upon the network locations of the computer **910** and the printer **916**. The combination of these segments forms a unique network path. However, a message moving from portable data terminal **926** to portable data terminal **922** may be routed along two different network paths. While both network paths include access point **906**, wired LAN backbone **908** and access point **902**, one network path includes first channel cell **930** while the other network path includes second channel cell **932**. Thus, depending upon system conditions and the system configuration, the message is routed via one of the two network paths. Such conditions may include cell traffic, required data rates and other factors.

FIG. 10 is a diagram illustrating a communication system **1000** built and operating according to the present invention wherein one of the access points routes communication between two portable terminal units operating on different channels within its cell. In the communication system **1000**, both a first access point **1002** and a second access point **1004** include both first and second channel radios. The first access point **1002** generates a first channel cell **1006** and a second channel cell **1008** within which portable data terminals **1012** and **1014** operate. Further, second access point **1004** generates a first channel cell **1010** and a second channel cell **1011** within which portable data terminals **1012** and **1014** may operate. In the embodiment, the system **1000** prefers to route communication on the first channel due to its characteristics although the portable data terminals may operate on either channel.

As illustrated, portable data terminal **1012** resides within both the first channel cell **1006** and the second channel cell

US 6,665,536 B1

13

1008 generated by access point 1002. However, portable data terminal 1014 resides only within the second channel cell 1008 of the access point 1002. Thus, in the transmission of a message from portable data terminal 1012 to portable data terminal 1014, access point 1002 receives the message from portable data terminal on the first channel radio and transmits the message to portable data terminal 1014 on the second channel. With reference to FIG. 6, the processing unit 612 receives the message via the first channel radio 616 across the bus interface 610. The processing unit 612 determines the destination of the message, and routes the message back across the bus interface 610 to the second channel radio 608 that transmits the message to portable data terminal 1014. Access point 1004 also provides multiple channel routing of messages between portable data terminals 1016 and 1018.

Without the multiple channel communication capabilities of the communication system 1000, an additional access point 1020 having a first channel cell 1022 would be required to facilitate communication with portable data terminals 1014 and 1018. The cost of such an additional access point 1020 would not only include the cost of the access point 1020 itself but the expense of connecting the access point 1020 to the wired LAN backbone 908 and AC power. The cost of such addition would far exceed the cost of the second channel radios in access points 1002 and 1004. Furthermore, in some installation, extensions of the wired LAN backbone 908 are not possible. Even if such access point 1020 were installed, the exemplary communication would require routing of messages between portable data terminal 1012 and 1014 across the wired LAN backbone 908. Such additional loading slows operation of the wired LAN backbone 908 and decreases system performance.

FIG. 11 is a block diagram illustrating an embodiment of a communication system 1100 according to the present invention wherein an access point 1102 uses a dedicated control/busy channel transmitter 1114 operating on a busy/control channel to manage transmissions between the access point 1100 and a plurality of roaming portable data terminals 1104 and 1106 within its cell. The communication system may also contain wired communication to a wired Ethernet backbone LAN 908.

The access device 1102 includes control circuitry 1120, a data transceiver 1118, a busy/control transmitter 1114 and antennas 1115 and 1117. The data transceiver 1117 supports communication on a communication channel (first channel) between the access point 1102 and wireless network devices operating within range of the access point, such as the portable data terminals 1104 and 1106. Further, the busy/control transmitter 1114 supports transmissions on the busy/control channel (second channel). The Ethernet transceiver 1115 supports communication between the backbone LAN 908 and the control circuitry 1120.

Portable data terminals 1104 and 1106 include terminal circuitry 1112, a data transceiver 1108 that communicates on the communication channel via antenna 1109 and a busy/control receiver 1110 that receives busy/control information via antenna 1111. As previously described, the communication channel and the busy/control channel are non-convergent and may operate concurrently in a single area or location. However, the access point 1102 must operate so as not to interfere with incoming transmissions by concurrently initiating a transmission. Thus, in one embodiment, transmissions on the communication channel and the control/busy channel are synchronized to prevent such conflicts.

The access point 1102 employs the busy/control transmitter 1114 to control operations within the first wireless

14

network cell. In one embodiment, the access point 1102 periodically transmits control parameters that the portable data terminals 1104 and 1106 use to synchronize with communications on the communication channel. For example, with the data transceiver of the communication channel operating in a spread spectrum mode, the busy/control transmitter 1114 transmits code spreading sequences, frequency hopping parameters and other operating parameters that allow the portable data terminals 1104 and 1106 to communicate within the cell on the communication channel. Such control information may be intermittently transmitted by the access point 1102 or may be continuously transmitted.

Additionally, the access point 1102 transmits a busy signal on the busy/control transmitter 1114 to authorize communication within the cell. To prevent portable data terminal 1104, for example, from transmitting while portable data terminal 1106 is communicating with the access point 1102, the access point 1102 transmits a busy signal on the busy/control channel using the busy/control transmitter 1114. The portable data terminal 1104 receives the busy signal and does not transmit information while such busy signal is active, perhaps entering a sleep mode instead and waking up periodically to determine availability. The busy signal may include a continuous transmission or periodic transmission. However, in both embodiments, portable data terminals 1104 and 1106 listen with their respective control/busy receivers 1110 prior to initiating communication with the access point 1102. Thus, upon roaming into range of the wireless access device 1102, the portable data terminals 1104 do not interfere with ongoing communication.

FIG. 12 is a drawing illustrating advantageous operation of the access device and portable data terminals of FIG. 11 when two roaming terminals encounter hidden terminal conditions. In particular, each of the portable data terminals 1208 and 1212 is configured to listen on the busy/control channel and to communicate on the communication channel only when the communication channel is clear (available). In this configuration, when no desire to communicate is present, the portable data terminals 1208 and 1212 need only occasionally check the busy/control channel to identify any outstanding messages or communication requests as transmitted by the access device 1202. If either portable data terminal 1208 or 1212 desires to participate on the communication channel (to initiate communication or to respond to awaiting messages or communication requests), that terminal need only monitor the busy/control channel long enough to identify that the communication channel is clear before responding to a poll on the communication channel. As before, the wireless access device 1202 may also periodically identify the communication channel mode and associated parameters as selected and reselected by the wireless access device 1202.

To fully appreciate this process, first assume that the portable data terminals 1208 and 1212 are not within range of the wireless access device 1202. Upon wandering within range of the access device 1202, the portable data terminal 1212 begins listening for transmissions on a busy/control channel. Within some time period thereafter, the access device 1202 participates on the busy/control channel to transmit current channel conditions and optionally, the currently selected communication channel definition (i.e., mode and parameters) and/or pending message and communication request indicators. After identifying a need to participate, the portable data terminal 1212 awaits a transmission from access device 1202 (on the busy/control channel) that the selected communication channel is clear (not in use). When the channel is clear, the portable data terminal 1212 begins participating thereon.

US 6,665,536 B1

15

Second, assume that, while the portable data terminal 1212 is engaged in ongoing communication with a computing device 1206 on a backbone LAN 908 via the access device 1202, the portable data terminal 1208 comes within range of the access device 1202 and desires to participate on the communication channel. The portable data terminal 1208 adapts itself to participate on the busy/control channel and identifies, in periodic transmissions from the access device 1202, that the communication channel is busy. Thus, the portable data terminal 1208 must monitor the busy/control channel to identify when the communication channel is clear before participating on the communication channel.

This operation works whether or not the portable data terminals 1208 and 1212 are within range of each other. In particular, portable data terminal 1208, portable data terminal 1212 and access device 1202 have transmission ranges illustrated by dashed circles 1210, 1214 and 1204, respectively. Although both portable data terminals 1208 and 1212 are within range of the access device 1202, neither are in range of each other and, thus, are referred to as "hidden" from each other. The access device 1202 is within range of both of the portable data terminals 1208 and 1212. If the portable data terminal 1208 attempted to transmit on the communication channel while the portable data terminal 1212 was transmitting, a collision would occur at the wireless access device 1202. However, this is not the case because both of the portable data terminals 1208 and 1212 must receive a communication channel clear indication on the busy/control channel from the access device 1202 that is in range of both, avoiding the hidden terminal problem. When participation is completed on the communication channel, the portable data terminals 1208 and 1212 resume monitoring of the busy/control channel.

Participation by the access device 1202 on the busy/control channel need only be by transmitting, although receiving might also be employed in case the busy/control channel is to be shared. Similarly, participation by the portable data terminals 1208 and 1212 need only be by receiving transmissions, although transmitting might also be employed. In particular, transmission might be employed by a wireless terminal on the busy/control channel if the wireless terminal does not support the currently selected communication channel, i.e., does not support the mode and associated parameters.

FIG. 13 is a block diagram illustrating an alternate embodiment of the communication system 1300 of the present invention wherein an access point 1302 includes a dedicated control/busy channel transceiver 1310 and roaming data terminals 1304 communicate with the access point 1302 using either frequency nimble multi-channel transceivers 1305 or a dedicated control/busy channel transceiver 1326. Thus, the communication system 1300 facilitates bi-directional communication on the busy/control channel so that the access point 1302 may optimize operation of the system 1300.

In addition to the busy/control transceiver 1310 coupled to antenna 1314, the access point includes control circuitry 1306, a data transceiver 1312 coupled to antenna 1316 that facilitates wireless communication on the communication channel and an Ethernet transceiver 1308 that couples the access point 1302 to the backbone LAN 908. Portable data terminal 1304 includes terminal circuitry 1112 and a multi-mode/multi-channel transceiver 1305 that allows the portable data terminal 1304 to communicate both on the busy/control channel and the communication channel. Portable data terminal 1320 includes terminal circuitry 1322, a busy/control transceiver 1324 coupled to antenna 1330 that allows

16

the portable data terminal 1320 to communicate on the busy/control channel and a data transceiver 1326 coupled to antenna 1328 that allows the portable data terminal to communicate on the communication channel.

Having separate radio units and antennas, the access device 1302 participates on: 1) a selected communication channel, servicing data exchanges in the communication network cell; and 2) the busy/control channel defined by predetermined mode and parameter information known to all wireless transmitters, controlling access to the selected communication channel. Such participation is often simultaneous, preventing a portable data terminal 1304 or 1320 from having to wait long on the busy/control channel for a transmission. Within a predefined maximum time period, the portable data terminal 1304 or 1320 receives transmissions from the access device 1302 identifying currently selected communication channel mode and associated parameters, should such be required. The access device 1302 periodically broadcasts such information on the busy/control channel to capture terminals that happen to need communication channel definitions (e.g., selected mode and parameters) to participate. The portable data terminal 1304 utilizes the identified mode and associated parameter information to switch the multi-mode transceiver 1305 over to the selected communication channel and begins participation thereon. Portable data terminal 1320 may also alter the operation of the data transceiver 1326 based upon the receipt from the access point 1302.

In operation, the wireless terminal 1304 participates on the busy/control channel except when it has a need to gain access to the selected communication channel. Thus, its operation in the system 1300 is satisfactory. By including only the terminal circuitry 1112 and one radio, the portable data terminal is less costly than the multi-radio portable data terminal 1320. However, because the wireless terminal 1320 includes two radios, the portable data terminal 1320 may place the data transceiver 1326 in a low power state, and only power up the busy/control channel transceiver 1324 to check in. Thus, portable data terminal 1320 may consume less power than portable data terminal 1304.

FIG. 14a is a block diagram illustrating a communication system 1400 according to the present invention wherein an access point 1402 and portable data terminals 1404 and 1406 operate on a deterministic first channel and a non-deterministic second channel and the system 1400 routes communications on the channels based upon system conditions and/or the requirements of a particular communication. To carry out such functionality, the access device 1402 may comprise control circuitry 1408, a wired LAN transceiver 1410 and either a single, configurable transceiver (for operating on both the deterministic and non-deterministic channels, not shown) or a single transceiver 1412 coupled to antenna 1413 for operating on the deterministic channel and a single transceiver 1414 coupled to antenna 1415 for operating on the non-deterministic channel.

In one embodiment, the deterministic channel allocates a particular communication bandwidth to each wireless device requiring communication, perhaps in a polled, token passing or time slotted implementation. Such operation may be required where many wireless devices reside within a single cell and compete for communication with the access point 1402. In the embodiment, the access point 1402 also allows all devices within the cell to compete for available bandwidth on the non-deterministic channel. However, the access point 1402 may provide overrides to dynamically reallocate bandwidth in the deterministic channel and to assign bandwidth on the non-deterministic channel as may be required for the particular operating conditions.

US 6,665,536 B1

17

With a single multi-mode transceiver 1418 coupled to antenna 1419 controlled by terminal circuitry 1416, the portable data terminal 1404 operates on either the deterministic channel or the non-deterministic channel at any time. Alternatively, portable data terminal 1406 having terminal circuitry 1420, a deterministic transceiver coupled to antenna 1425 and a non-deterministic transceiver 1422 coupled to antenna 1423 communicate on both the deterministic channel and non-deterministic channel simultaneously.

Independent of their differing constructions, the portable data terminals 1404 and 1406 may determine which channel to operate upon. During data transfer operations wherein data transfer rates are not critical, portable data terminal 1404 may determine that the deterministic channel provides sufficient bandwidth. In that case, the portable data terminal 1404 configures its multi-mode transceiver 1418 to operate on the deterministic channel. However, during voice message transfer operations, the portable data terminal 1404 may determine that the bandwidth of the deterministic channel is not satisfactory. In that case, the terminal circuitry 1404 would configure the multi-mode transceiver 1418 to operate on the non-deterministic channel.

FIG. 14b is a diagram illustrating a communication system 1450 according to the present invention that facilitates both wired and wireless communications. The communication system 1450 includes a wired backbone 1452 and at least one access point 1456 that supports communication over a deterministic, time bounded first wireless channel and a non-deterministic, contention access second wireless channel. Along with the access point 1456, the communication system 1450 may also include a PBX (Private Broadcast Exchange) system 1454, one or more of a computer 1454, and other typical wired network devices interconnected by the wired backbone 1452. Additionally, the communication system 1450 comprises a plurality of wireless network devices, such as wireless terminals 1470, 1472 and 1474, which may be portable hand-held devices, mobile computing devices, laptop computers, wireless peripherals, etc.

Communication upon the wired backbone 1452 may be accomplished according to various communication techniques. In one embodiment, communication upon the wired backbone 1452 occurs via an STM (Synchronous Transfer Mode) protocol wherein the wired backbone 1452 serves as an STM backbone. With the STM protocol, a particular bandwidth is provided for each communication link established between a sending and a receiving device attached to the wired backbone 1452.

In another embodiment, communication upon the wired backbone 1452 is carried out using an ATM (Asynchronous Transfer Mode) protocol in which bandwidth between a sending and a receiving device on the wired backbone 1452 is adjusted based upon immediate communication requirements. In such operation, the wired backbone 1452 serves as an ATM backbone. Operation according to the ATM protocol allows for variations in data transmission bandwidths as is immediately required but that provides an average bandwidth over time.

The at least one access point 1456 provides a link between the wireless and wired communications within the communication system 1450. The access point 1456 includes a time bounded adapter 1458 connected to an antenna 1460 which provides wireless communication on the deterministic, time bounded first wireless channel governed by a first wireless protocol. The access point 1456 also includes a contention adapter 1462 connected to an antenna 1464 which provides

18

wireless communications on the non-deterministic, contention access second wireless channel governed by a second wireless protocol. Alternatively, part or all of the circuitry underlying the adapters 1458 and 1462 may be combined into a single unit to share common underlying functionality.

The wireless network device 1470 includes either a dual purpose transceiver or two transceivers for communicating on the first and second wireless channels via the first and second wireless protocols, respectively. A transceiver in the wireless network device 1472 only supports communication on the second wireless channel pursuant to the second wireless protocol. Likewise, a transceiver in the wireless network device 1474 supports communication on the first wireless channel pursuant to the first wireless protocol. For example, the wireless network device 1474 might comprise a portable phone unit operating using, e.g., PCS (Personal Communication Service) or other telephony protocol as the first wireless protocol.

Although direct communication is possible, to manage the first wireless channel, the at least one access point 1456 relays wireless communication between the wireless network devices 1470 and 1474 if both participate on the first wireless channel. If the at least one access point 1456 is the only access point involved that services the two devices 1470 and 1474, such relaying need not involve the wired protocol on the backbone 1452. If the device 1470 happens to communicate via the second wireless channel, the access point 1456 internally translates and relays communications between the devices 1470 and 1474. Similarly, if the device 1470 intends to communicate with a wired network device using either the first or second wireless protocol, the access point 1456 utilizes the first or second wireless protocols, respectively, to communicate with the device 1470. The access point 1456 also communicates with the target wired network device, e.g., the computer 1454, via the wired communication protocol. Relaying between the wired and wireless channels also requires translation.

If more than one access point is coupled to the wired backbone 1452, for example, to support many more wireless network devices, roaming wireless network devices and/or extended coverage regions, the access points only utilize wired backbone bandwidth if necessary. Each access point attempts to minimize external bandwidth (of wired and wireless channels) by preferring internally performed relaying and, when needed, translation (between the first and second wireless protocols, or between the wired protocol and either the first or the second wireless protocol).

The PBX system 1454 connects the wired backbone 1452 through a switched telephone network to other communication systems such as the system 1452. This facilitates communications between all wireless and wired network devices, such as the computer 1454, the device 1470 and remote network devices (devices) connected elsewhere to the switched telephone network. In circuit switched applications, a VLAN (virtual local area network) can be established between wired and wireless network devices coupled to the wired backbone 1452. Such coupling also includes remote network devices coupled via the PBX system 1454.

FIG. 15 is a diagram illustrating the access point 1402 and portable data terminals 1404 and 1406 of FIG. 14a wherein a system 1500 routes various transmissions according to system conditions such as channel activity, data type and data priority. In the system 1500, access point 1402 forms cell 1512 while access point 1504, operating on only a single channel, forms cell 1514. Thus, while access point 1402

US 6,665,536 B1

19

must determine how to allocate wireless communications among the deterministic channel and non-deterministic channel in its cell 1512, access point 1504 routes all communications on its only channel.

In a first example of the operation of the system 1500, data from computer system 1502 is transmitted to portable data terminal 1406. The computer system 1502 transmits the data through the wired LAN backbone 908 to the access point 1402. The access point 1402, having a deterministic transceiver 1412 and a non-deterministic transceiver 1414, routes the data through one of the transceivers. Based upon the type of data, the quantity of data, the rate at which data may be passed on either channel, the amount of traffic on the channels and other conditions, the control circuitry 1408 in the access point 1402 routes the data on either the non-deterministic channel via the non-deterministic transceiver 1414 or on the deterministic channel via the deterministic transceiver 1412. In the present example, the data to be transferred has a relatively low priority and the access point routes the data on the deterministic channel to the portable data terminal 1404. Likewise, print jobs from the computer 1402 to the printer 1510 would also have relatively low priority and be transmitted via the non-deterministic channel.

Next, consider data transmissions from scanner 1508 to computer system 1502. During operation, the scanner transmits an image to the computer system 1502 for decoding and the computer system 1502 returns decoded information at which point the scanner ceases scanning. Rapid transmission between the scanner 1508 and the computer system 1502 reduces the time within which the scanner 1508 performs scans. Thus, rapid transmissions may reduce energy consumption in the scanning process that drains battery life of the scanner 1508. Thus, the scanner 1508 and access point 1402 both attempt to transmit data on the non-deterministic channel at a relatively higher data transfer rate. However, if the non-deterministic channel is unavailable, the transmission on the deterministic channel may be satisfactory.

In the case of a voice message transmission from portable data terminal 1506 in cell 1514 to portable data terminal 1404 in cell 1512, transmission on the deterministic channel may be unsatisfactory. Thus, upon an incoming voice message transmission, the access point 1402 may reallocate the deterministic channel allocating additional bandwidth for the voice message. In an alternative operation, the access point 1402 may interrupt communication on the non-deterministic channel and transmit the voice message to the portable terminal unit 1404 on the non-deterministic channel. Thus, in the mode of operation of the system 1500 modifies its operation to provide sufficient bandwidth for the voice message.

In view of the above detailed description of the present invention and associated drawings, other modifications and variations will now become apparent to those skilled in the art. It should also be apparent that such other modifications and variations may be effected without departing from the spirit and scope of the present invention as set forth in the claims that follow.

What is claimed is:

1. An access point for communicatively coupling a first roaming wireless device and a second roaming wireless device to a wired link, the access point comprising:
 - a housing;
 - a control circuit disposed in the housing;
 - a wired transceiver, disposed in the housing, that is communicatively coupled to the control circuit and the wired link;

20

a first wireless transceiver, disposed in the housing, that is communicatively coupled to the control circuit, the first wireless transceiver operating on a first wireless communication channel to communicatively couple with the first roaming wireless device;

a second wireless transceiver, disposed in the housing, that is communicatively coupled to the control circuit, the second wireless transceiver operating on a second wireless communication channel to communicatively couple with the second roaming wireless device; and the control circuit accommodates communications between the first wireless transceiver and the second wireless transceiver exclusive of the wired link.

2. The access point of claim 1, further comprising a bus interface communicatively coupling the control circuit to the first and second wireless transceivers and the wired transceiver.

3. The access point of claim 2, wherein the bus interface is substantially compliant with a bus standard.

4. The access point of claim 3, wherein the bus standard is the PCI standard.

5. The access point of claim 1, wherein the wired transceiver accommodates communication with an ethernet network.

6. The access point of claim 1, wherein the wired transceiver accommodates communication with a token-ring network.

7. The access point of claim 1, wherein the wired transceiver accommodates communication with an asynchronous transfer mode network.

8. The access point of claim 1, wherein the wired transceiver accommodates communication with a packetized network.

9. The access point of claim 1, wherein the first wireless transceiver supports a substantially non-deterministic media access protocol and the second wireless transceiver supports a substantially deterministic media access protocol.

10. The access point of claim 1, wherein the first wireless transceiver and the second wireless transceiver support substantially distinct non-deterministic media access protocols.

11. The access point of claim 1, wherein the first wireless transceiver and the second wireless transceiver operate independently to form a first communication cell and a second communication cell.

12. The access point of claim 1, wherein the control circuit synchronizes transmissions on the first radio channel and the second radio channel to minimize conflicts between transmissions on one wireless transceiver and receipts on the other wireless transceiver.

13. The access point of claim 1, wherein the wired link is a local area network.

14. The communication network of claim 10, wherein the first wireless communication channel is a radio frequency (RF) channel.

15. An access point for establishing communications with a wired link, the access point comprising:

a first wireless transceiver operating to establish a first wireless cell;

a second wireless transceiver operating to establish a second wireless cell;

the first and second wireless transceivers operating such that the first and second cells are substantially overlapping;

a control circuit that communicatively couples the first and second wireless transceivers to one another;

US 6,665,536 B1

21

a wired transceiver that communicatively couples the control circuit to the wired link; and

the control circuit communicatively couples the first wireless transceiver and the wired transceiver.

16. The access point of claim 15, wherein the first and second wireless transceivers each comprise processing circuitry that supports a communication protocol.

17. The access point of claim 15, wherein the control circuit allows communications between the first wireless transceiver and the second wireless transceiver exclusive of the wired link.

18. The access point of claim 15, wherein the first wireless transceiver supports a substantially non-deterministic media access protocol and the second wireless transceiver supports a substantially deterministic media access protocol.

19. The access point of claim 15, wherein the first wireless transceiver and the second wireless transceiver support substantially distinct non-deterministic media access protocols.

20. A communication network comprising:

- a wired LAN;
- a plurality of access points coupled via the wired LAN, each of the plurality of access points comprising:
 - a housing;
 - a control circuit disposed in the housing;
 - wired transceiver, disposed in the housing, that is communicatively coupled to the control circuit and the wired LAN;
 - a first wireless transceiver, disposed in the housing, that is communicatively coupled to the control circuit and operates on a first wireless communication channel;
 - a second wireless transceiver, disposed in the housing, that is communicatively coupled to the control circuit and operates on a second wireless communication channel; and
 - the control circuit accommodates communications between the first wireless transceiver and the second wireless transceiver exclusive of the wired LAN;
- a first roaming wireless device comprising a third wireless transceiver that operates on the first wireless communication channel; and
- a second roaming wireless device comprising a fourth wireless transceiver that operates on the second wireless communication channel.

21. An access point for establishing communications with a wired link, the access point comprising:

- processing circuitry operating to send and receive data according to a first protocol; and
- interface circuitry operable to:
 - receive data from the processing circuitry according to the first protocol;
 - send data to a plurality of wireless transceivers operating on independent wireless communication channels, according to at least a second protocol independent of the first protocol;
 - send data to a wired transceiver operating on the wired link, according to a third protocol independent of the first and second protocols;
 - receive data from the plurality of wireless transceivers according to at least the second protocol independent of the first protocol;
 - receive data from the wired transceiver according to the third protocol independent of the first and second protocols; and

22

send data to the processing circuitry according to the first protocol.

22. The communication network of claim 20, wherein the first roaming device operates only on the first wireless communication channel.

23. The communication network of claim 20, wherein the first roaming wireless device and the second roaming wireless device have different transmission characteristics.

24. The communication network of claim 20, wherein the first roaming wireless device and the second roaming wireless device incorporate different data throughput capabilities.

25. The communication network of claim 20, wherein the first roaming wireless device and the second roaming wireless device operate independently to form a first communication cell and a second communication cell, respectively.

26. The communication network of claim 20, wherein the radius of the first communication cell substantially equals the radius of the second communication cell.

27. The communication network of claim 20, wherein the wired transceiver accommodates communication with an Ethernet network.

28. The communication network of claim 20, wherein the wired transceiver accommodates communication with a token-ring network.

29. The communication network of claim 20, wherein the wired transceiver accommodates communication with an asynchronous transfer mode network.

30. The communication network of claim 20, wherein the wired transceiver accommodates communication with a packetized network.

31. The communication network of claim 20, wherein the first wireless transceiver supports a substantially non-deterministic media access protocol and the second wireless transceiver supports a substantially deterministic media access protocol.

32. The communication network of claim 20, wherein the first wireless transceiver and the second wireless transceiver support substantially distinct non-deterministic media access protocols.

33. The communication network of claim 20, wherein the third wireless transceiver is a PCMCIA card.

34. A communication system, comprising:

- a wired LAN;
- a plurality of access points coupled via the wired LAN, each of the plurality of access points comprising:
 - a housing;
 - a control circuit disposed in the housing;
 - a wired transceiver, disposed in the housing, that is configurable to communicatively couple the control circuit to the wired LAN;
 - a first wireless transceiver, disposed in the housing, that is communicatively coupled to the control circuit, the first wireless transceiver operating pursuant to a substantially deterministic, time bounded wireless communication protocol; and
 - a second wireless transceiver, disposed in the housing, that is communicatively coupled to the control circuit, the second wireless transceiver operating pursuant to a substantially non-deterministic contention access wireless communication protocol; and
- a plurality of roaming wireless devices that each wirelessly communicate with at least one of the first and second wireless transceivers.

35. The access point of claim 21, wherein the second and third protocols are the same and comply with PCI bus standards.

US 6,665,536 B1

23

36. The access point of claim 21, wherein the processing circuitry is programmed with a network configuration to selectively route data through the interface circuitry to the plurality of wireless transceivers and the wired link.

37. The access point of claim 21, further comprising at least one acceptor for modularly receiving the plurality of wireless transceivers.

38. The access point of claim 37, wherein the plurality of transceivers are carried by at least one PCMCIA card.

39. The access point of claim 21, wherein the plurality of wireless transceivers operate independently to form a plurality of communication cells.

40. The access point of claim 39, wherein the plurality of communication cells are formed by the plurality of wireless transceivers operating at different data rates.

41. The access point of claim 39, wherein the plurality of communication cells are formed by the plurality of wireless transceivers operating at different power levels.

42. The access point of claim 21, wherein the independent wireless communication channels are differentiated by a characteristic selected from the group consisting of frequencies, modulation schemes and code spreading schemes.

43. An access point for establishing communications with a wired link, the access point comprising:

a housing;

a PCMCIA interface capable of modularly receiving into the housing a plurality of wireless transceivers operating on independent wireless communication channels; a wired transceiver in the housing operating on the wired link;

interface circuitry in the housing operable to communicate with wireless transceivers modularly received via the PCMCIA interface and with the wired transceiver; and

processing circuitry in the housing coupled to the interface circuitry to control communications by the wireless transceivers modularly received via the PCMCIA interface and by the wired transceiver.

44. The access point of claim 43, wherein the interface circuitry comprises a PCI bus interface for communicating with the wireless transceivers modularly received via the PCMCIA interface and with the wired transceiver according to PCI bus standards.

45. The access point of claim 43, wherein the processing circuitry is programmed with a network configuration to selectively route data through the interface circuitry to the plurality of wireless transceivers and the wired link.

46. The access point of claim 43, wherein the plurality of wireless transceivers operate independently to form a plurality of communication cells.

47. The access point of claim 46, wherein the plurality of communication cells are formed by the plurality of wireless transceivers operating at different data rates.

48. The access point of claim 46, wherein the plurality of communication cells are formed by the plurality of wireless transceivers operating at different power levels.

49. A communication system, comprising:

a wired LAN;

a plurality of access point coupled via the wired LAN, each of the plurality of access points comprising:

a housing;

a control circuit disposed in the housing;

a wired transceiver, disposed in the housing, that is configurable to communicatively couple the control circuit to a wired local area network; and

24

a wireless transceiver system, disposed in the housing, that is communicatively coupled to the control circuit, the wireless transceiver system contemporaneously operating on first and second communication channels; and

a plurality of roaming wireless devices that each wirelessly communicate with the wireless transceiver system using at least one of the first and second communication channels.

50. An access point for establishing communications with a wired link, the access point comprising:

a housing;

an interface system for modularly receiving into the housing a plurality of wireless transceivers operating on independent wireless communication channels;

interface circuitry in the housing operable to communicate with wireless transceivers modularly received via the interface system; and

processing circuitry in the housing coupled to the interface circuitry to control communications effected by wireless transceivers modularly received via the interface system.

51. An access point for establishing communications with a wired link, the access point comprising:

a housing;

receiving means for modularly receiving into the housing a plurality of wireless transceivers operating on independent wireless communication channels;

interface means in the housing for communicating with wireless transceivers modularly received by the receiving means; and

processing means in the housing coupled to the interface means for controlling communications by wireless transceivers modularly received by the receiving means.

52. A method of establishing communications with a wired link through an access point, the method comprising:

modularly receiving at least one wireless transceiver in the access point, the at least one wireless transceiver being selected from a plurality of wireless transceivers operating on independent wireless communication channels; and

communicating data and control information with the at least one wireless transceiver modularly received in the access point according to at least one protocol selected from a plurality of protocols supported by the plurality of wireless transceivers operating on independent wireless communication channels.

53. The access point of claim 52, wherein the plurality of wireless transceivers carried by the plurality of cards have substantially different operating characteristics.

54. An access point for establishing communications with a wired link, the access point comprising:

a housing;

receiving means for modularly receiving into the housing a plurality of wireless transceivers for operating on independent wireless communication channels;

US 6,665,536 B1

25

interface means in the housing for communicating with wireless transceivers modularly received by the receiving means; and

processing means in the housing coupled to the interface means for controlling communications by wireless transceivers modularly received by the receiving means.

55. A method of establishing communications with a wired link through an access point, the method comprising: modularly receiving at least one wireless transceiver in the access point, the at least one wireless transceiver

26

being selected from a plurality of wireless transceivers operating on independent wireless communication channels; and

communicating data and control information with the at least one wireless transceiver modularly received in the access point according to at least one protocol selected from a plurality of protocols supported by the plurality of wireless transceivers operating on independent wireless communication channels.

* * * * *

EXHIBIT 13



US006421714B1

(12) **United States Patent**
Rai et al.

(10) **Patent No.:** **US 6,421,714 B1**
(45) **Date of Patent:** **Jul. 16, 2002**

(54) **EFFICIENT MOBILITY MANAGEMENT
SCHEME FOR A WIRELESS INTERNET
ACCESS SYSTEM**

(75) Inventors: **Girish Rai, Bartlett; Philip M.
Parsons, Lisle, both of IL (US); Mool
Chuah, Eatontown, NJ (US)**

(73) Assignee: **Lucent Technologies, Murray Hill, NJ
(US)**

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/138,681**

(22) Filed: **Aug. 24, 1998**

Related U.S. Application Data

(60) Provisional application No. 60/061,915, filed on Oct. 14,
1997.

(51) Int. Cl.⁷ **G06F 13/00**

(52) U.S. Cl. **709/217; 709/225; 709/238;
709/250; 455/422; 455/432**

(58) Field of Search **709/217, 218,
709/220, 223, 225, 227, 229, 230, 238,
250, 249, 313; 455/422, 432, 433, 435**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,475,735 A * 12/1995 Williams et al. 455/403

5,796,729 A * 8/1998 Greaney et al. 370/345
5,958,016 A * 9/1999 Chang et al. 709/229
6,055,441 A * 4/2000 Wieand et al. 455/557
6,061,346 A * 5/2000 Nordman 370/352
6,061,650 A * 5/2000 Malkin et al. 704/228
6,091,952 A * 7/2000 Xu et al. 455/433
6,151,628 A * 11/2000 Xu et al. 709/225

* cited by examiner

Primary Examiner—Viet D. Vu

(74) *Attorney, Agent, or Firm*—Morgan & Finnegan

(57) **ABSTRACT**

A wireless data network that provides communications with a Point-to-Point Protocol (PPP) server is disclosed. A home network includes a home mobile switching center and a wireless end system, the home mobile switching center including a home registration server and a home inter-working function, the wireless end system including an end registration agent, the end registration agent being coupled to the home registration server. The wireless data network also includes a PPP server, wherein a message is coupleable from the end system through the home inter-working function to the PPP server.

28 Claims, 31 Drawing Sheets




FIG. 1
(PRIOR ART)




FIG. 2





FIG. 3



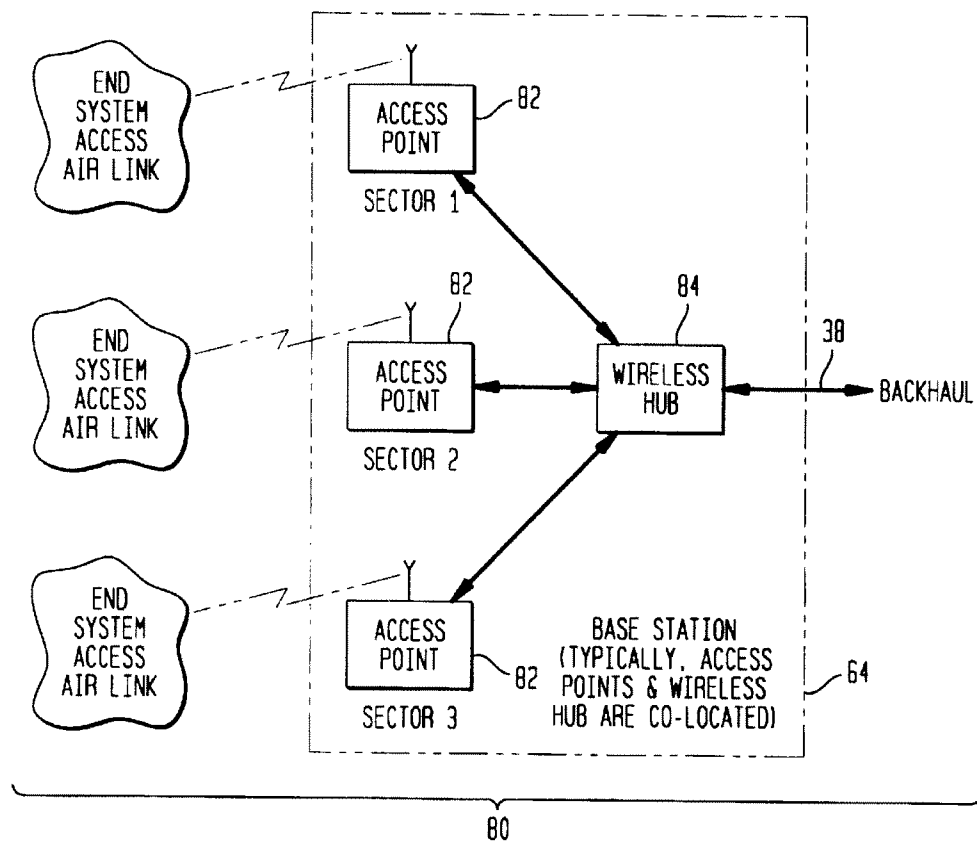
U.S. Patent

Jul. 16, 2002

Sheet 4 of 31

US 6,421,714 B1

FIG. 4




U.S. Patent

Jul. 16, 2002

Sheet 5 of 31

US 6,421,714 B1

FIG. 5




U.S. Patent

Jul. 16, 2002

Sheet 6 of 31

US 6,421,714 B1

FIG. 6



U.S. Patent

Jul. 16, 2002

Sheet 7 of 31

US 6,421,714 B1

FIG. 7




FIG. 8





FIG. 9



U.S. Patent

Jul. 16, 2002

Sheet 8 of 31

US 6,421,714 B1

FIG. 10

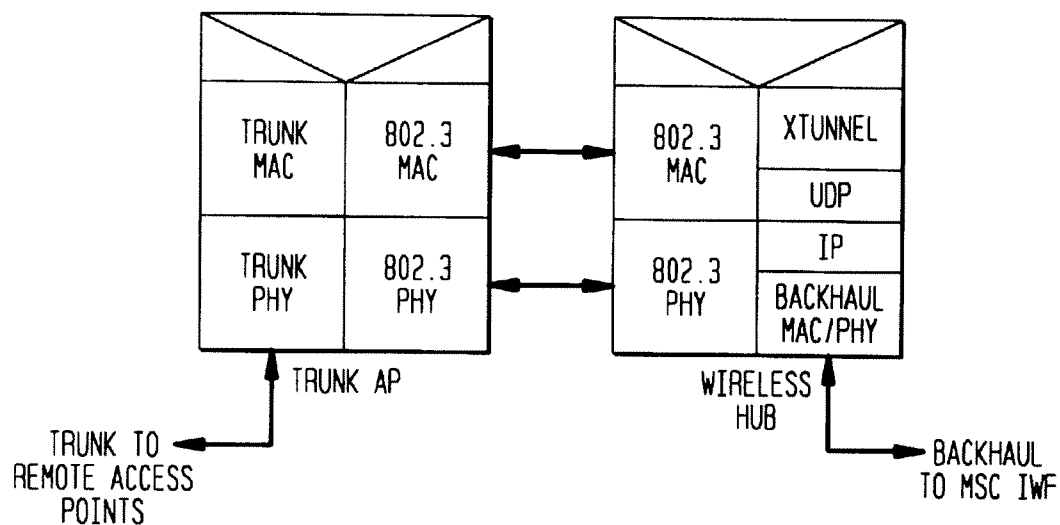



FIG. 11




U.S. Patent

Jul. 16, 2002

Sheet 9 of 31

US 6,421,714 B1

FIG. 12



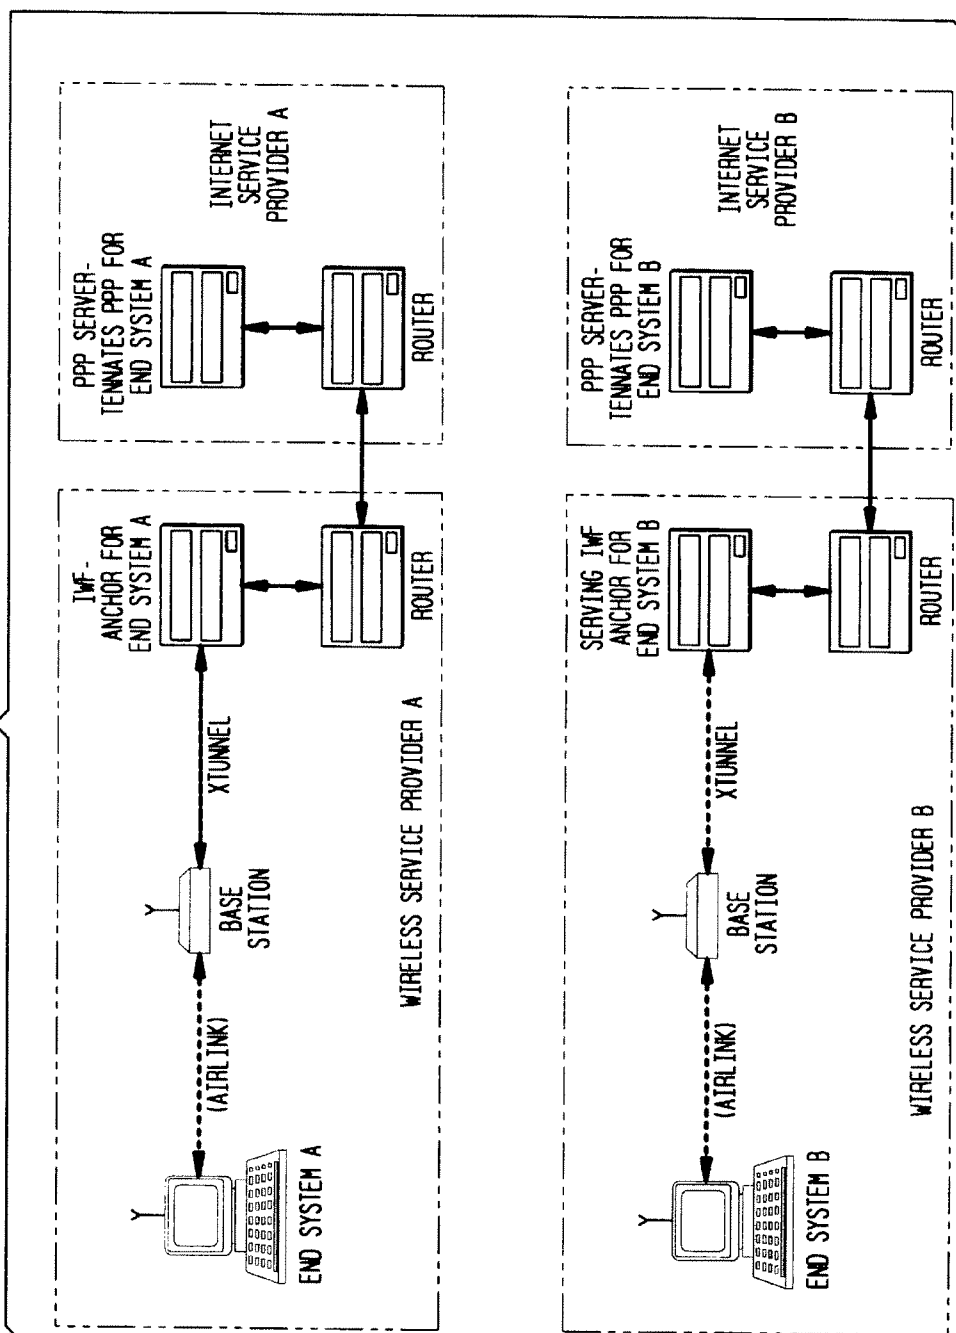
U.S. Patent

Jul. 16, 2002

Sheet 10 of 31

US 6,421,714 B1

FIG. 13



U.S. Patent

Jul. 16, 2002

Sheet 11 of 31

US 6,421,714 B1

FIG. 14

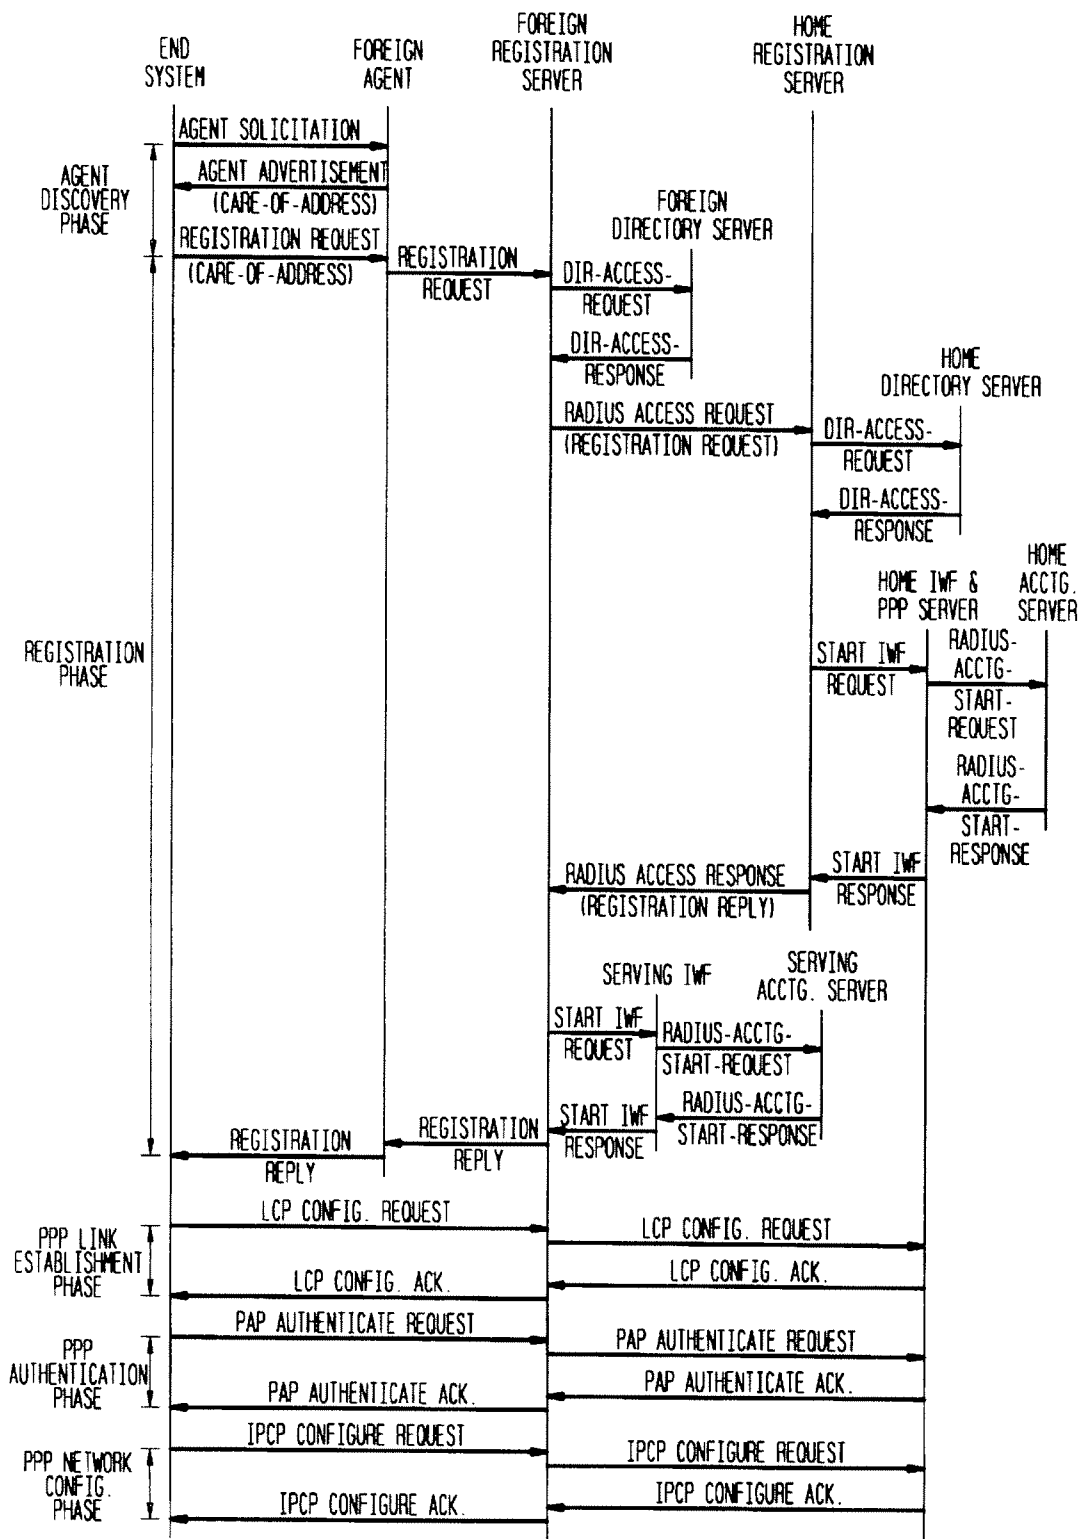


FIG. 15





FIG. 16



U.S. Patent

Jul. 16, 2002

Sheet 14 of 31

US 6,421,714 B1

FIG. 17





FIG. 18




U.S. Patent

Jul. 16, 2002

Sheet 15 of 31

US 6,421,714 B1

FIG. 19



U.S. Patent

Jul. 16, 2002

Sheet 16 of 31

US 6,421,714 B1

FIG. 20





FIG. 21




U.S. Patent

Jul. 16, 2002

Sheet 17 of 31

US 6,421,714 B1

FIG. 22



U.S. Patent

Jul. 16, 2002

Sheet 18 of 31

US 6,421,714 B1

FIG. 23

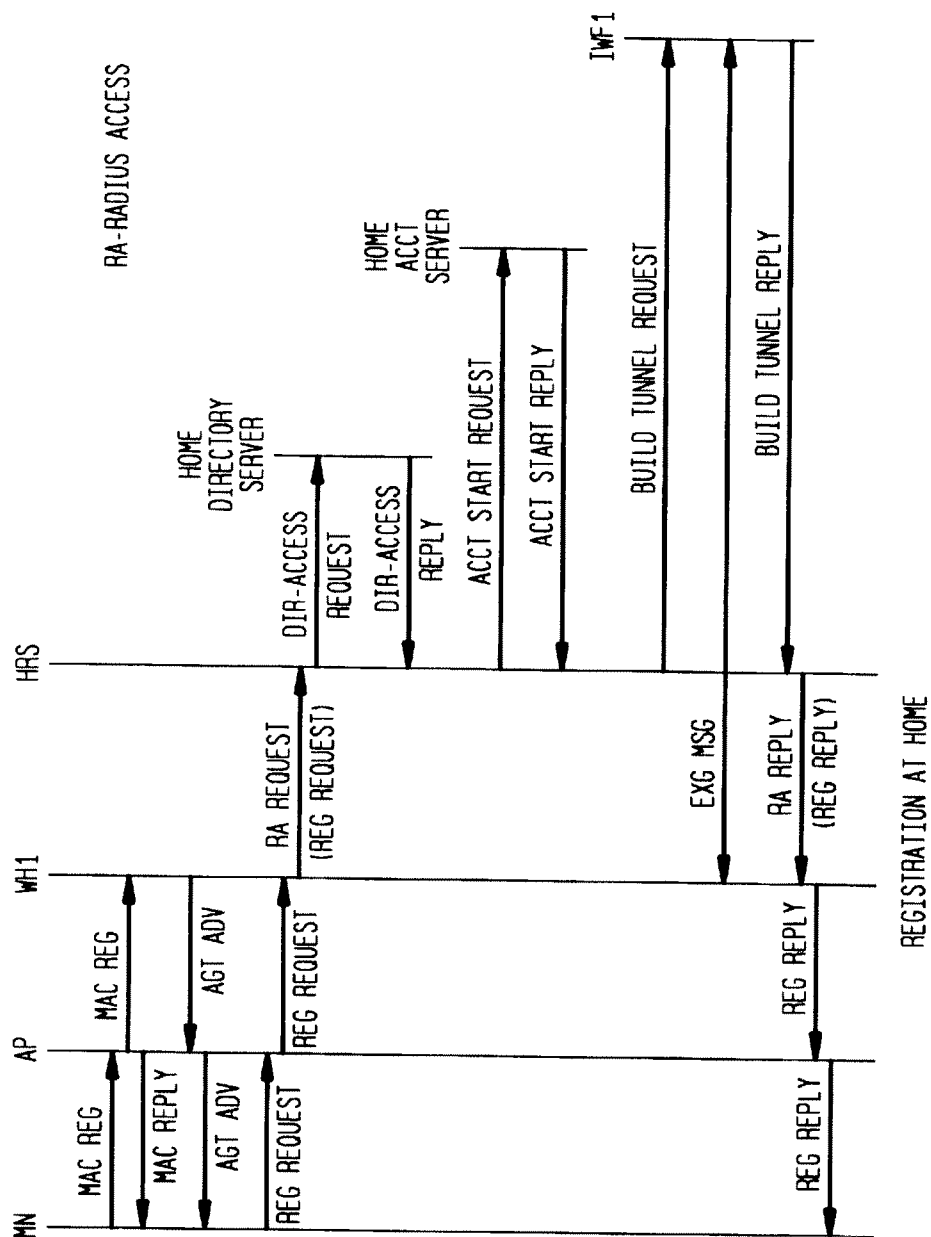
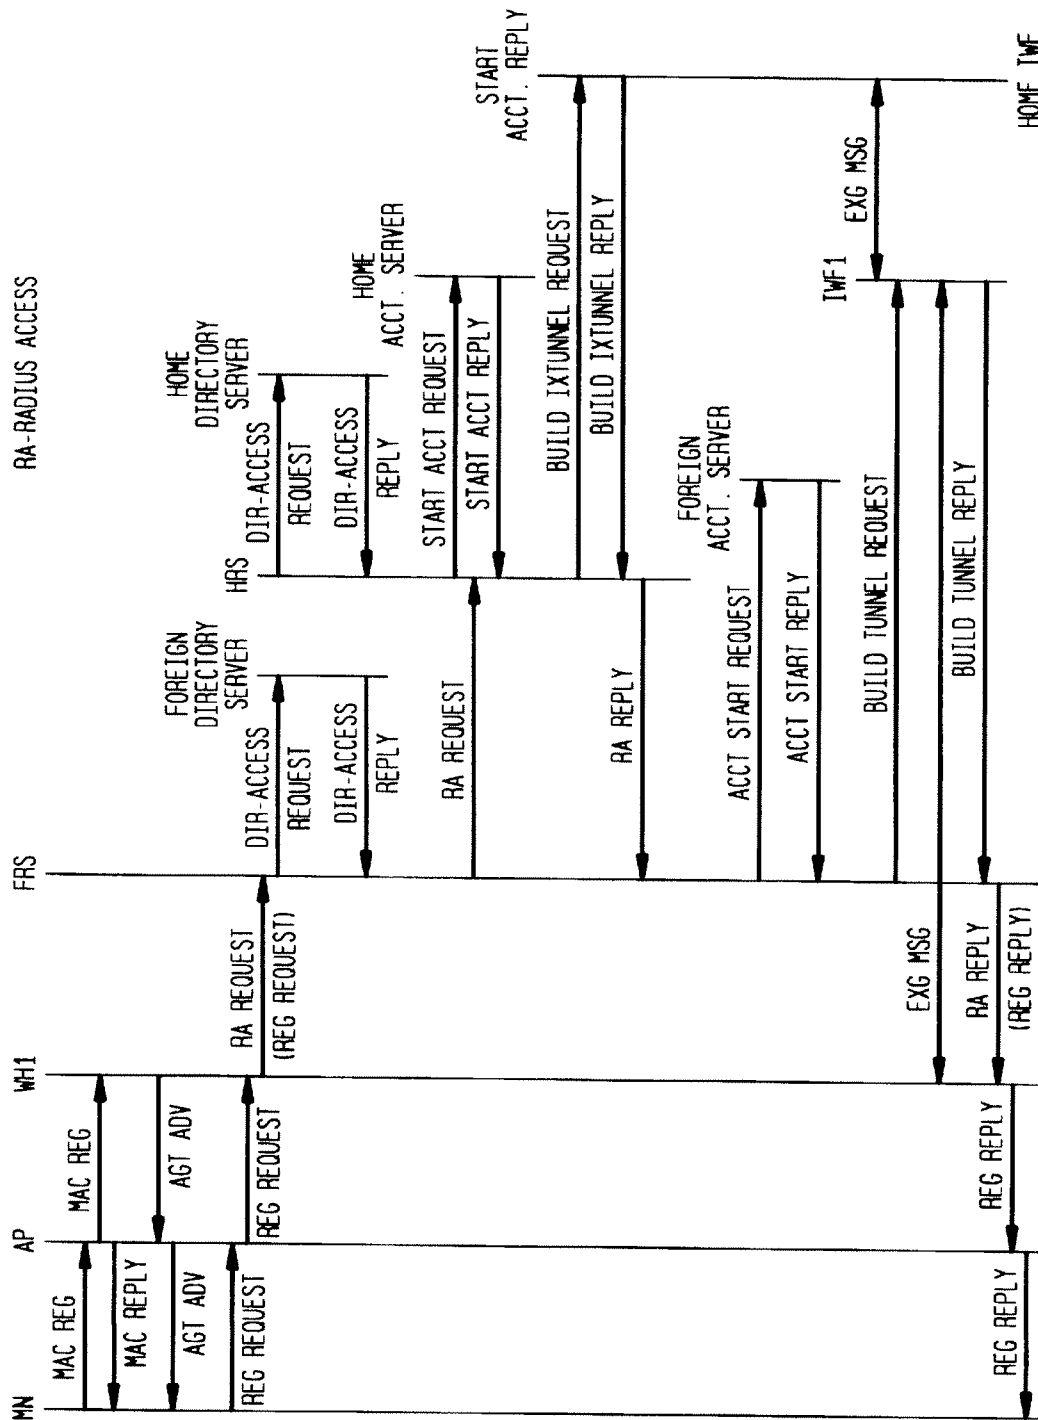


FIG. 24




U.S. Patent

Jul. 16, 2002

Sheet 20 of 31

US 6,421,714 B1

FIG. 25




U.S. Patent

Jul. 16, 2002

Sheet 21 of 31

US 6,421,714 B1

FIG. 26




U.S. Patent

Jul. 16, 2002

Sheet 22 of 31

US 6,421,714 B1

FIG. 27



PROTOCOL STACKS FOR ROAMING MOBILE WHERE HOME IWF IS ALSO A PPP SERVER


U.S. Patent

Jul. 16, 2002

Sheet 23 of 31

US 6,421,714 B1

FIG. 28



PROTOCOL STACKS FOR ROAMING MOBILE WHERE HOME IMF AND PPP SERVER ARE SEPARATE


U.S. Patent

Jul. 16, 2002

Sheet 24 of 31

US 6,421,714 B1

FIG. 29



U.S. Patent

Jul. 16, 2002

Sheet 25 of 31

US 6,421,714 B1

FIG. 30





FIG. 31




U.S. Patent

Jul. 16, 2002

Sheet 26 of 31

US 6,421,714 B1

FIG. 32




U.S. Patent

Jul. 16, 2002

Sheet 27 of 31

US 6,421,714 B1

FIG. 33



LOCAL HANDOFF

U.S. Patent

Jul. 16, 2002

Sheet 28 of 31

US 6,421,714 B1

FIG. 34

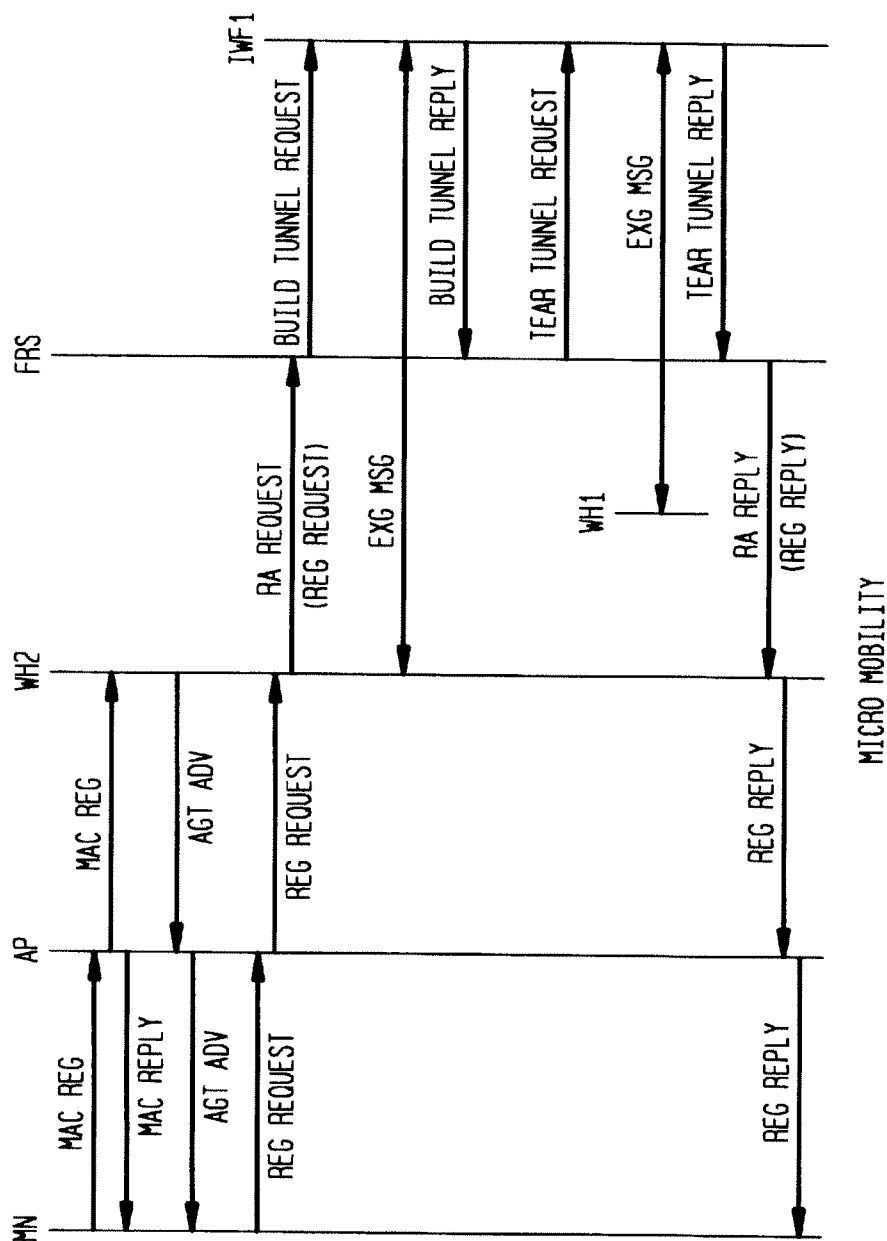


FIG. 35

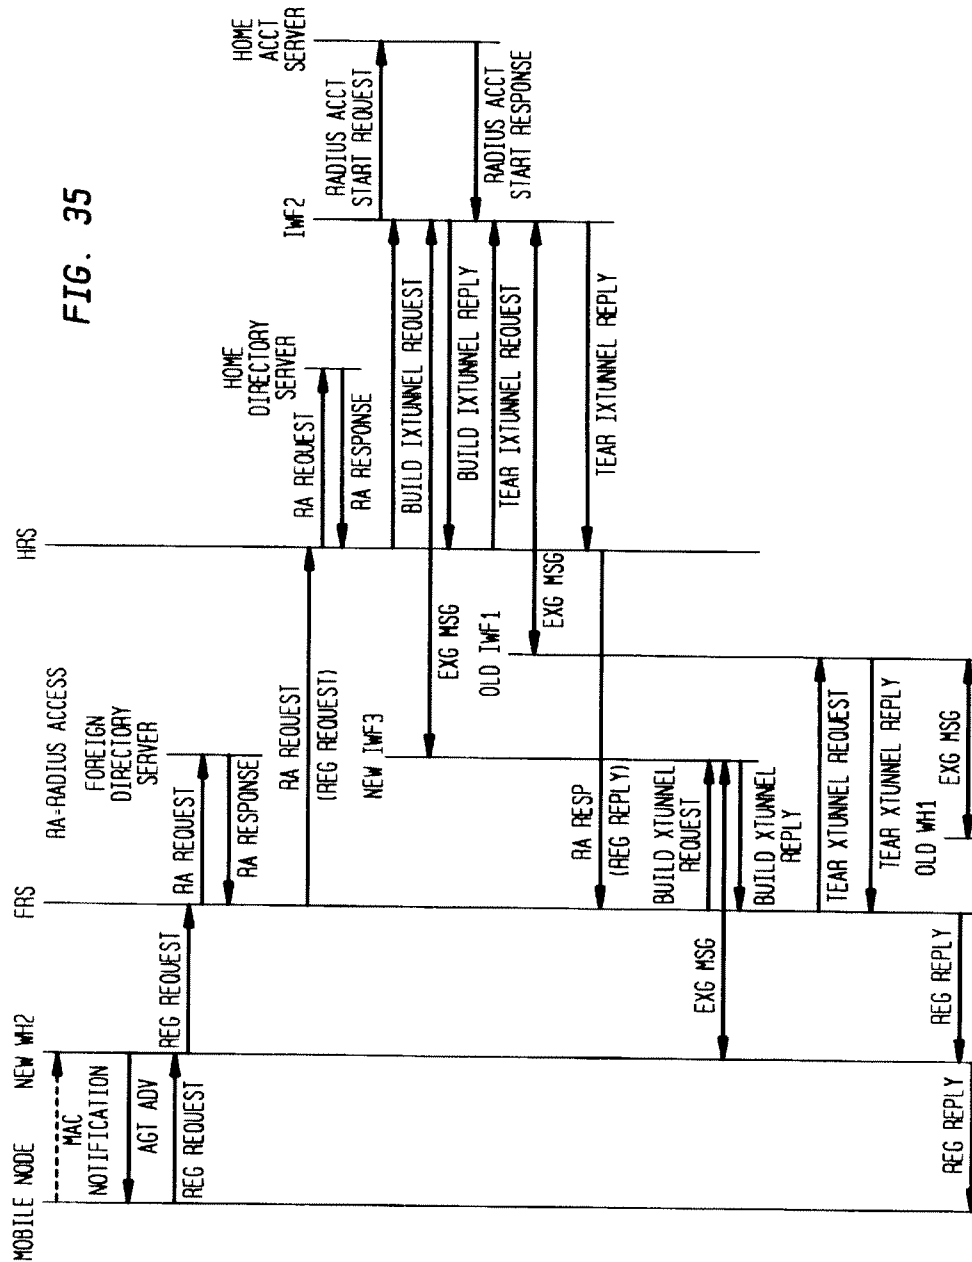


FIG. 36

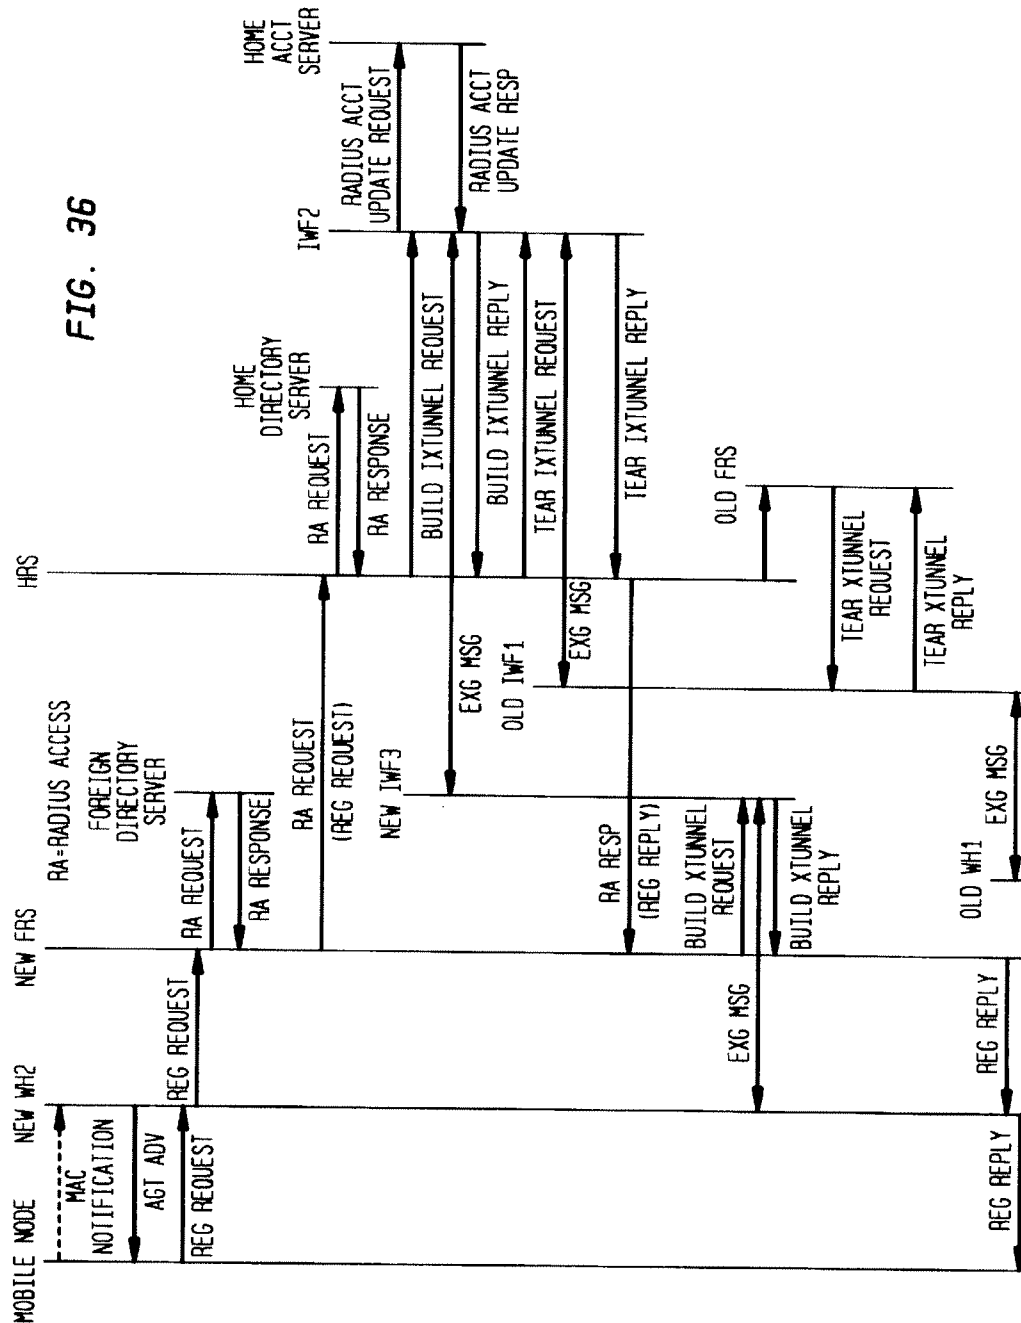



FIG. 37



US 6,421,714 B1

1

EFFICIENT MOBILITY MANAGEMENT SCHEME FOR A WIRELESS INTERNET ACCESS SYSTEM

Priority benefit of the Oct. 14, 1997 filing date of 5
provisional application Ser. No. 60/061,915 is hereby
claimed.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to a wireless data network,
and more particularly to an efficient mobility management
scheme for a wireless Internet access system.

2. Description of Related Art

FIG. 1 depicts three business entities, whose equipment,
working together typically provide remote internet access to
user computers 2 through user modems 4. User computers 2
and modems 4 constitute end systems.

The first business entity is the telephone company (telco)
that owns and operates the dial-up plain old telephone
system (POTS) or integrated services data network (ISDN)
network. The telco provides the media in the form of public
switched telephone network (PSTN) 6 over which bits (or
packets) can flow between users and the other two business
entities.

The second business entity is the internet service provider
(ISP). The ISP deploys and manages one or more points of
presence (POPs) 8 in its service area to which end users
connect for network service. An ISP typically establishes a
POP in each major local calling area in which the ISP
expects to subscribe customers. The POP converts message
traffic from the PSTN run by the telco into a digital form to
be carried over intranet backbone 10 owned by the ISP or
leased from an intranet backbone provider like MCI, Inc. An
ISP typically leases fractional or full T1 lines or fractional or
full T3 lines from the telco for connectivity to the PSTN.
The POPs and the ISP's medium data center 14 are con-
nected together over the intranet backbone through router
12A. The data center houses the ISP's web servers, mail
servers, accounting and registration servers, enabling the
ISP to provide web content, e-mail and web hosting services
to end users. Future value added services may be added by
deploying additional types of servers in the data center. The
ISP also maintains router 12A to connect to public internet
backbone 20. In the current model for remote access, end
users have service relationships with their telco and their ISP
and usually get separate bills from both. End users access the
ISP, and through the ISP, public internet 20, by dialing the
nearest POP and running a communication protocol known
as the Internet Engineering Task Force (IETF) point-to-point
protocol (PPP).

The third business entity is the private corporation which
owns and operates its own private intranet 18 through router
12B for business reasons. Corporate employees may
access corporate network 18 (e.g., from home or while on
the road) by making POTS/ISDN calls to corporate remote
access server 16 and running the IETF PPP protocol. For
corporate access, end users only pay for the cost of con-
necting to corporate remote access server 16. The ISP is not
involved. The private corporation maintains router 12B to
connect an end user to either corporate intranet 18 or public
internet 20 or both.

End users pay the telco for the cost of making phone calls
and for the cost of a phone line into their home. End users
also pay the ISP for accessing the ISP's network and

2

services. The present invention will benefit wireless service
providers like Sprint PCS, PrimeCo, etc. and benefit internet
service providers like AOL, AT&T Worldnet, etc.

Today, internet service providers offer internet access
services, web content services, e-mail services, content
hosting services and roaming to end users. Because of low
margins and no scope of doing market segmentation based
on features and price, ISPs are looking for value added
services to improve margins. In the short term, equipment
vendors will be able to offer solutions to ISPs to enable them
to offer faster access, virtual private networking (which is
the ability to use public networks securely as private net-
works and to connect to intranets), roaming consortiums,
push technologies and quality of service. In the longer term,
voice over internet and mobility will also be offered. ISPs
will use these value added services to escape from the low
margin straitjacket. Many of these value added services fall
in the category of network services and can be offered only
through the network infrastructure equipment. Others fall in
the category of application services which require support
from the network infrastructure, while others do not require
any support from the network infrastructure. Services like
faster access, virtual private networking, roaming, mobility,
voice, quality of service, quality of service based accounting
all need enhanced network infrastructure. The invention
described here will be either directly provide these enhanced
services or provide hooks so that these services can be added
later as future enhancements. Wireless service providers will
be able to capture a larger share of the revenue stream. The
ISP will be able to offer more services and with better market
segmentation.

SUMMARY OF THE INVENTION

The present invention provide end users with remote
wireless access to the public internet, private intranets and
internet service providers. Wireless access is provided
through base stations in a home network and base stations in
foreign networks with interchange agreements.

It is an object of the present invention to provide a
wireless packet switched data network for end users that
divides mobility management into local, micro, macro and
global connection handover categories and minimizes hand-
off updates according to the handover category. It is another
object to integrate MAC handoff messages with network
handoff messages. It is a further object of the present
invention to separately direct registration functions to a
registration server and direct routing functions to inter-
working function units. It is yet another object to provide an
intermediate XTunnel channel between a wireless hub (also
called access hub AH) and an inter-working function unit
(IWF unit) in a foreign network. It is yet another object to
provide an IXTunnel channel between an inter-working
function unit in a foreign network and an inter-working
function unit in a home network. It is yet another object to
enhance the layer two tunneling protocol (L2TP) to support
a mobile end system. It is yet another object to perform
network layer registration before the start of a PPP commu-
nication session.

According to one embodiment of the invention, a wireless
data network which provides communications with a PPP
server is disclosed. A home network includes a home mobile
switching center and a wireless end system, the home mobile
switching center including a home registration server and a
home inter-working function, the wireless end system
including an end registration agent, the end registration
agent being coupled to the home registration server. The

US 6,421,714 B1

3

wireless data network also includes a PPP server, wherein a message is coupleable from the end system through the home inter-working function to the PPP server.

According to another embodiment of the invention, a data network to communicate with a first PPP protocol module is disclosed. The data network comprises a mobile end system and a home inter-working function. The mobile end system is operable in first and second modes, the first mode providing internet access services, the second mode providing remote intranet access services, the mobile end system including a second PPP protocol module, PPP data frames being transportable between the first and second PPP protocol modules. The home inter-working function is incorporated in the first PPP protocol module when the mobile end system operates in the first mode, the home inter-working function being coupled to the first PPP protocol module operating externally when the mobile end system operates in the second mode.

BRIEF DESCRIPTION OF DRAWINGS

The invention will be described in detail in the following description of preferred embodiments with reference to the following figures wherein:

FIG. 1 is a configuration diagram of a known remote access architecture through a public switched telephone network;

FIG. 2 is a configuration diagram of a remote access architecture through a wireless packet switched data network according to the present invention;

FIG. 3 is a configuration diagram of selected parts of the architecture of the network of FIG. 2 showing a roaming scenario;

FIG. 4 is a configuration diagram of a base station with local access points;

FIG. 5 is a configuration diagram of a base station with remote access points;

FIG. 6 is a configuration diagram of a base station with remote access points, some of which are connected using a wireless trunk connection;

FIG. 7 is a diagram of a protocol stack for a local access point;

FIG. 8 is a diagram of a protocol stack for a remote access point with a wireless trunk;

FIG. 9 is a diagram of a protocol stack for a relay function in the base station for supporting remote access points with wireless trunks;

FIG. 10 is a diagram of protocol stacks for implementing the relay function depicted in FIG. 9;

FIG. 11 is a diagram of protocol stacks for a relay function in the base station for supporting local access points;

FIG. 12 is a configuration diagram of selected parts of the architecture of the network of FIG. 2 showing a first end system registering in the home network from the home network and a second system registering in the home network from a foreign network using a home inter-working function for an anchor;

FIG. 13 is a configuration diagram of selected parts of the architecture of the network of FIG. 2 showing a first end system registering in the home network from the home network and a second system registering in the home network from a foreign network using a serving inter-working function for an anchor;

FIG. 14 is a ladder diagram of the request and response messages to register in a home network from a foreign network and to establish, authenticate and configure a data link;

4

FIG. 15 is a configuration diagram of selected parts of the architecture of the network of FIG. 2 showing registration requests and responses for registering a mobile in a home network from the home network;

FIG. 16 is a configuration diagram of selected parts of the architecture of the network of FIG. 2 showing registration requests and responses for registering a mobile in a home network from a foreign network;

FIG. 17 is a configuration diagram of protocol stacks showing communications between an end system in a home network and an inter-working function in the home network where the cell site has local access points;

FIG. 18 is a configuration diagram of protocol stacks showing communications between an end system in a home network and an inter-working function in the home network where the cell site has remote access points coupled to a wireless hub through a wireless trunk;

FIG. 19 is a configuration diagram of protocol stacks showing communications between a base station coupled to a roaming end system and a home inter-working function;

FIG. 20 is a configuration diagram of protocol stacks showing communications between an end system in a home network through an inter-working function in the home network to an internet service provider;

FIG. 21 is a configuration diagram of protocol stacks showing communications between an end system in a foreign network and a home registration server in a home network during the registration phase;

FIG. 22 is a processing flow diagram showing the processing of accounting data through to the customer billing system;

FIGS. 23 and 24 are ladder diagrams depicting the registration process for an end system in a home network and in a foreign network, respectively;

FIGS. 25 and 26 are protocol stack diagrams depicting an end system connection in a home network where a PPP protocol terminates in an inter-working function of the home network and where the PPP protocol terminates in an ISP or intranet, respectively;

FIGS. 27 and 28 are protocol stack diagrams depicting an end system connection in a foreign network where a PPP protocol terminates in an inter-working function of the foreign network and where the PPP protocol terminates in an ISP or intranet, respectively;

FIG. 29 illustrates end systems connected via ethernet to a wireless modem where PPP protocol is encapsulated in an ethernet frame;

FIG. 30 illustrates an ethernet frame format;

FIG. 31 illustrates XWD Header fields;

FIG. 32 illustrates end systems connected via a local area network to a wireless router where PPP protocol terminates at the wireless router;

FIGS. 33, 34 and 35 are ladder diagrams depicting a local handoff scenario, a micro handoff scenario and a macro handoff scenario, respectively;

FIG. 36 is a ladder diagram depicting a global handoff scenario where the foreign registration server changes and where home inter-working function does not change; and

FIG. 37 is a ladder diagram depicting a global handoff scenario where both the foreign registration server and the home inter-working function change.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

The present invention provides computer users with remote access to the internet and to private intranets using

US 6,421,714 B1

5

virtual private network services over a high speed, packet switched, wireless data link. These users are able to access the public internet, private intranets and their internet service providers over a wireless link. The network supports roaming, that is, the ability to access the internet and private intranets using virtual private network services from anywhere that the services offered by the present system are available. The network also supports handoffs, that is, the ability to change the point of attachment of the user to the network without disturbing the PPP link between the PPP client and the PPP server. The network targets users running horizontal internet and intranet applications. These applications include electronic mail, file transfer, browser based WWW access and other business applications built around the internet. Because the network will be based on the IETF standards, it is possible to run streaming media protocols like RTP and conferencing protocols like H.323 over it.

Other internet remote access technologies that are already deployed or are in various stages of deployment include: wire line dial-up access based on POTS and ISDN, XDSL access, wireless circuit switched access based on GSM/CDMA/TDMA, wireless packet switched access based on GSM/CDMA/TDMA, cable modems, and satellite based systems. However, the present system offers a low cost of deployment, ease of maintenance, a broad feature set, scalability, an ability to degrade gracefully under heavy load conditions and support for enhanced network services like virtual private networking, roaming, mobility and quality of service to the relative benefit of users and service providers.

For wireless service providers who own personal communications system (PCS) spectrum, the present system will enable them to offer wireless packet switched data access services that can compete with services provided by the traditional wire line telcos who own and operate the PSTN. Wireless service providers may also decide to become internet service providers themselves, in which case, they will own and operate the whole network and provide end to end services to users.

For internet service providers the present system will allow them to by-pass the telcos (provided they purchase or lease the spectrum) and offer direct end to end services to users, perhaps saving access charges to the telcos, which may increase in the future as the internet grows to become even bigger than it is now.

The present systems flexible so that it can benefit wireless service providers who are not internet service providers and who just provide ISP, internet or private intranet access to end users. The system can also benefit service providers who provide wireless access and internet services to end users. The system can also benefit service providers who provide wireless access and internet services but also allow the wireless portion of the network to be used for access to other ISPs or to private intranets.

In FIG. 2, end systems 32 (e.g., based on, for example, Win 95 personal computer) connect to wireless network 30 using external or internal modems. These modems allow end systems to send and receive medium access control (MAC) frames over air link 34. External modems attach to the PC via a wired or wireless link. External modems are fixed, and, for example, co-located with roof top mounted directional antennae. External modems may be connected to the user's PC using any one of following means: 802.3, universal serial bus, parallel port, infra-red, or even an ISM radio link. Internal modems are preferably PCMCIA cards for laptops and are plugged into the laptop's backplane. Using a small

6

omni-directional antenna, they send and receive MAC frames over the air link. End systems can also be laptops with a directional antenna, a fixed wireless station in a home with a directional antenna connected via AC lines, and other alternatives. Wide-area wireless coverage is provided by base stations 36. The base station 36 can employ a 5-channel reuse communication scheme as described in U.S. patent application Ser. No. 08/998,505, filed on Dec. 26, 1997. The range of coverage provided by base stations 36 depends on factors like link budget, capacity and coverage. Base stations are typically installed in cell sites by PCS (personal communication services) wireless service providers. Base stations multiplex end system traffic from their coverage area to the system's mobile switching center (MSC) 40 over wire line or microwave backhaul network 38.

The system is independent of the MAC and PHY (physical) layer of the air link and the type of modem. The architecture is also independent of the physical layer and topology of backhaul network 38. The only requirements for the backhaul network are that it must be capable of routing internet protocol (IP) packets between base stations and the MSC with adequate performance. At Mobile Switching Center 40 (MSC 40), packet data inter-working function (IWF) 52 terminates the wireless protocols for this network. IP router 42 connects MSC 40 to public internet 44, private intranets 46 or to internet service providers 46. Accounting and directory servers 48 in MSC 40 store accounting data and directory information. Element management server 50 manages the equipment which includes the base stations, the IWFs and accounting/directory servers.

The accounting server will collect accounting data on behalf of users and send the data to the service provider's billing system. The interface supported by the accounting server will send accounting information in American Management Association (AMA) billing record format, or any other suitable billing format, over a TCP/IP (transport control protocol/internet protocol) transport to the billing system (which is not shown in the figure).

The network infrastructure provides PPP (point-to-point protocol) service to end systems. The network provides (1) fixed wireless access with roaming (log-in anywhere that the wireless coverage is available) to end systems and (2) low speed mobility and hand-offs. When an end system logs on to a network, in it may request either fixed service (i.e., stationary and not requiring handoff services) or mobile service (i.e., needing handoff services). An end system that does not specify fixed or mobile is regarded as specifying mobile service. The actual registration of the end system is the result of a negotiation with a home registration server based on requested level of service, the level of services subscribed to by the user of the end system and the facilities available in the network.

If the end system negotiates a fixed service registration (i.e., not requiring handoff services) and the end system is located in the home network, an IWF (inter-working function) is implemented in the base station to relay traffic between the end user and a communications server such as a PPP server (i.e., the point with which to be connected, for example, an ISP PPP server or a corporate intranet PPP server or a PPP server operated by the wireless service provider to provide customers with direct access to the public internet). It is anticipated that perhaps 80% of the message traffic will be of this category, and thus, this architecture distributes IWF processing into the base stations and avoids message traffic congestion in a central mobile switching center.

If the end system requests mobile service (from a home network or a foreign network) or if the end system request

US 6,421,714 B1

7

roaming service (i.e., service from the home network through a foreign network), two IWFs are established: a serving IWF typically established in the base station of the network to which the end system is attached (be it the home network or a foreign network) and a home IWF typically established in mobile switching center MSC of the home network. Since this situation is anticipated to involve only about 20% of the message traffic, the message traffic congestion around the mobile switching center is minimized. The serving IWF and the wireless hub may be co-located in the same nest of computers or may even be programmed in the same computer so that a tunnel using an XTunnel protocol need not be established between the wireless hub and the serving IWF.

However, based on available facilities and the type and quality of service requested, a serving IWF in a foreign network may alternatively be chosen from facilities in the foreign MSC. Generally, the home IWF becomes an anchor point that is not changed during the communications session, while the serving IWF may change if the end system moves sufficiently.

The base station includes an access hub and at least one access point (be it remote or collocated with the access hub). Typically, the access hub serves multiple access points. While the end system may be attached to an access point by a wire or cable according to the teachings of this invention, in a preferred embodiment the end system is attached to the access point by a wireless "air link", in which case the access hub is conveniently referred to as a wireless hub. While the access hub is referred to as a "wireless hub" throughout the description herein, it will be appreciated that an end system coupled through an access point to an access hub by wire or cable is an equivalent implementation and is contemplated by the term "access hub".

In the invention, an end system includes an end user registration agent (e.g., software running on a computer of the end system, its modem or both) that communicates with an access point, and through the access point to a wireless hub. The wireless hub includes a proxy registration agent (e.g., software running on a processor in the wireless hub) acting as a proxy for the end user registration agent. Similar concepts used in, for example, the IETF proposed Mobile IP standard are commonly referred to as a foreign agent (FA). For this reason, the proxy registration agent of the present system will be referred to as a foreign agent, and aspects of the foreign agent of the present system that differ from the foreign agent of Mobile IP are as described throughout this description.

Using the proxy registration agent (i.e., foreign agent FA) in a base station, the user registration agent of an end system is able to discover a point of attachment to the network and register with a registration server in the MSC (mobile switching center) of the home network. The home registration server determines the availability of each of the plural inter-working function modules (IWFs) in the network (actually software modules that run on processors in both the MSC and the wireless hubs) and assigns IWF(s) to the registered end system. For each registered end system, a tunnel (using the XTunnel protocol) is created between the wireless hub in the base station and an inter-working function (IWF) in the mobile switching center (MSC), this tunnel transporting PPP frames between the end system and the IWF.

As used herein, the XTunnel protocol is a protocol that provides in-sequence transport of PPP data frames with flow control. This protocol may run over standard IP networks or

8

over point-to-point networks or over switched networks like ATM data networks or frame relay data networks. Such networks may be based on T1 or T3 links or based on radio links, whether land based or space based. The XTunnel protocol may be built by adapting algorithms from L2TP (level 2 transport protocol). In networks based on links where lost data packets may be encountered, a re-transmission feature may be a desirable option.

The end system's PPP peer (i.e., a communications server) may reside in the IWF or in a corporate intranet or ISP's network. When the PPP peer resides in the IWF, an end system is provided with direct internet access. When the PPP peer resides in an intranet or ISP, an end system is provided with intranet access or access to an ISP. In order to support intranet or ISP access, the IWF uses the layer two tunneling protocol (L2TP) to connect to the intranet or ISP's PPP server. From the point of view of the intranet or ISP's PPP server, the IWF looks like a network access server (NAS). PPP traffic between the end system and the IWF is relayed by the foreign agent in the base station.

In the reverse (up link) direction, PPP frames traveling from the end system to the IWF are sent over the MAC and air link to the base station. The base station relays these frames to the IWF in the MSC using the XTunnel protocol. The IWF delivers them to a PPP server for processing. For internet access, the PPP server may be in the same machine as the IWF. For ISP or intranet access, the PPP server is in a private network and the IWF uses the layer two tunneling protocol (L2TP) to connect to it.

In the forward (down link) direction, PPP frames from the PPP server are relayed by the IWF to the base station using the XTunnel protocol. The base station de-tunnels down link frames and relays them over the air link to the end system, where they are processed by the end system's PPP layer.

To support mobility, support for hand-offs are included. The MAC layer assists the mobility management software in the base station and the end system to perform hand-offs efficiently. Hand-offs are handled transparently from the peer PPP entities and the L2TP tunnel. If an end system moves from one base station to another, a new XTunnel is created between the new base station and the original IWF. The old XTunnel from the old base station will be deleted. PPP frames will transparently traverse the new path.

The network supports roaming (i.e., when the end user connects to its home wireless service provider through a foreign wireless service provider). Using this feature, end systems are able to roam away from the home network to a foreign network and still get service, provided of course that the foreign wireless service provider and the end system's home wireless service provider have a service agreement.

In FIG. 3, roaming end system 60 has traveled to a location at which foreign wireless service provider 62 provides coverage. However, roaming end system 60 has a subscriber relationship with home wireless service provider 70. In the present invention, home wireless service provider 70 has a contractual relationship with foreign wireless service provider 62 to provide access services. Therefore, roaming end system 60 connects to base station 64 of foreign wireless service provider 62 over the air link. Then, data is relayed from roaming end system 60 through base station 64, through serving IWF 66 of foreign wireless service provider 62, to home IWF 72 of home wireless service provider 70, or possibly through home IWF 72 of home wireless service provider 70 to internet service provider 74.

An inter-service provider interface, called the I-interface, is used for communications across wireless service provider

US 6,421,714 B1

9

(WSP) boundaries to support roaming. This interface is used for authenticating, registering and for transporting the end system's PPP frames between the foreign WSP and the home WSP.

PPP frames in the up link and the down link directions travel through the end system's home wireless service provider (WSP). Alternatively, PPP frames directly transit from the foreign WSP to the destination network. The base station in the foreign WSP is the end system's point of attachment in the foreign network. This base station sends (and receives) PPP frames to (and from) a serving IWF in the foreign WSP's mobile switching center. The serving IWF connects over the I-interface to the home IWF using a layer two tunnel to transport the end system's PPP frames in both directions. The serving IWF in the foreign WSP collects accounting data for auditing. The home IWF in the home WSP collects accounting data for billing.

The serving IWF in the foreign WSP may be combined with the base station in the same system, thus eliminating the need for the X-Tunnel.

During the registration phase, a registration server in the foreign WSP determines the identity of the roaming end system's home network. Using this information, the foreign registration server communicates with the home registration server to authenticate and register the end system. These registration messages flow over the I-interface. Once the end system has been authenticated and registered, a layer two tunnel is created between the base station and the serving IWF using the XTUNNEL protocol and another layer two tunnel is created between the serving IWF and the home IWF over the I-interface. The home IWF connects to the end system's PPP peer as before, using L2TP (level 2 tunnel protocol). During hand-offs, the location of the home IWF and the L2TP tunnel remains fixed. As the end system moves from one base station to another base station, a new tunnel is created between the new base station and the serving IWF and the old tunnel between the old base station and the serving IWF is deleted. If the end system moves far enough, so that a new serving IWF is needed, a new tunnel will be created between the new serving IWF and the home IWF. The old tunnel between the old serving and the home will be deleted.

To support roaming, the I-interface supports authentication, registration and data transport services across wireless service provider boundaries. Authentication and registration services are supported using the IETF Radius protocol. Data transport services to transfer PPP frames over a layer two tunnel are supported using the I-XTunnel protocol. This protocol is based on the IETF L2TP protocol.

As used in this description, the term home IWF refers to the IWF in the end system's home network. The term serving IWF refers to the IWF in the foreign network which is temporarily providing service to the end system. Similarly, the term home registration server refers to the registration server in the end system's home network and the term foreign registration server refers to the registration server in the foreign network through which the end system registers while it is roaming.

The network supports both fixed and dynamic IP address assignment for end systems. There are two types of IP addresses that need to be considered. The first is the identity of the end system in its home network. This may be a structured user name in the format user@domain. This is different from the home IP address used in mobile IP. The second address is the IP address assigned to the end system

10

via the PPP IPCP address negotiation process. The domain sub-field of the home address is used to identify the user's home domain and is a fully qualified domain name. The user sub-field of the home address is used to identify the user in the home domain. The User-Name is stored on the end system and in the subscriber data-base at the MSC and is assigned to the user when he or she subscribes to the service. The domain sub-field of the User-Name is used during roaming to identify roaming relationships and the home registration server for purposes of registration and authentication. Instead of the structured user name another unique identifier may be used to identify the user's home network and the user's identity in the home network. This identifier is sent in the registration request by the end system.

The PPP IPCP is used to negotiate the IP address for the end system. Using IP configuration protocol IPCP, the end system is able to negotiate a fixed or dynamic IP address.

Although the use of the structured user-name field and the non-use of an IP address as the home address is a feature that characterizes the present system over a known mobile IP, the network may be enhanced to also support end systems that have no user-name and only a non-null home address, if mobile IP and its use in conjunction with PPP end systems becomes popular. The PPP server may be configured by the service provider to assign IP addresses during the IPCP address assignment phase that are the same as the end system's home IP address. In this case, the home address and the IPCP assigned IP address will be identical.

In FIG. 4, base station 64 and air links from end systems form wireless sub-network 80 that includes the air links for end user access, at least one base station (e.g., station 64) and at least one backhaul network (e.g., 38 of FIG. 2) from the base station to MSC 40 (FIG. 2). The wireless sub-network architecture of, for example, a 3-sectored base station includes the following logical functions.

1. Access point function. Access points 82 perform MAC layer bridging and MAC layer association and disassociation procedures. An access point includes a processor (preferably in the form of custom application specific integrated circuit ASIC), a link to a wireless hub (preferably in the form of an Ethernet link on a card or built into the ASIC), a link to an antenna (preferably in the form of a card with a data modulator/demodulator and a transmitter/receiver), and the antenna to which the end system is coupled. The processor runs software to perform a data bridging function and various other functions in support of registration and mobility handovers as further described herein. See discussion with respect to FIGS. 7, 8 and 11.

2. Access points (APs) take MAC layer frames from the air link and relay them to a wireless hub and vice versa. The MAC layer association and disassociation procedures are used by APs to maintain a list of end system MAC addresses in their MAC address filter table. An AP will only perform MAC layer bridging on behalf of end systems whose MAC addresses are present in the table. An access point and its associated wireless hub are typically co-located. In its simplest form, an access point is just a port into a wireless hub. When the APs and the wireless hub are co-located in the same cell site, they may be connected together via a IEEE 802.3 link. Sometimes, access points are located remotely from the wireless hub and connected via a long distance link like a wired T1 trunk or even a wireless trunk. For multi-sector cells, multiple access points (i.e., one per sector) are used.

3. Wireless hub function. Wireless hub 84 performs the foreign agent (FA) procedures, backhaul load balancing

US 6,421,714 B1

11

(e.g., over multiple T1's), backhaul network interfacing, and the xtunnel procedures. When support for quality of service (QOS) is present, the wireless hub implements the support for QOS by running the xtunnel protocol over backhauls with different QOS attributes. In a multi-sector cell site, a single wireless hub function is typically shared by multiple access points.

A wireless hub includes a processor, a link to one or more access points (preferably in the form of an Ethernet link on a card or built into an ASIC), and a link to a backhaul line. The backhaul line is typically a T1 or T3 communications line that terminates in the mobile switching center of the wireless service provider. The link to the backhaul line formats data into a preferred format, for example, an Ethernet format, a frame relay format or an ATM format. The wireless hub processor runs software to support data bridging and various other functions as described herein. See discussion with respect to FIGS. 9, 10 and 11.

The base station design supports the following types of cell architectures.

1. Local AP architecture. In a local AP architecture, access points have a large (>2 km, typically) range. They are co-located in the cell site with the wireless hub (FIG. 4). Access points may be connected to the wireless hub using an IEEE 802.3 network or may be directly plugged into the wireless hub's backplane or connected to the wireless hub using some other mechanism (e.g. universal serial bus, printer port, infra-red, etc.). It will be assumed that the first alternative is used for the rest of this discussion. The cell site may be omni or sectorized by adding multiple access points and sectorized antennas to a wireless hub.
2. Remote AP architecture. In a remote AP architecture, access points usually have a very small range, typically around 1 km radius. They are located remotely (either indoors or outdoors) from the wireless hub. A T1 or a wireless trunk preferably links remote access points to the cell site where the wireless hub is located. From the cell site, a wire line backhaul or a microwave link is typically used to connect to the IWF in the MSC. If wireless trunking between the remote AP and the wireless hub is used, omni or sectorized wireless radios for trunking are utilized. The devices for trunking to remote access points are preferably co-located with the wireless hub and may be connected to it using an IEEE 802.3 network or may be directly plugged into the wireless hub's backplane. These devices will be referred to by the term trunk AP.
3. Mixed AP architecture. In a mixed architecture, the wireless sub-network will have to support remote and local access points. Remote access points may be added for hole filling and other capacity reasons. As described earlier, T1 or wireless trunks may be used to connect the remote AP to the wireless hub.

FIG. 5 shows a cell with three sectors using local APs only. The access points and the wireless hub are co-located in the base station and are connected to each other with 802.3 links.

FIG. 6 shows an architecture with remote access points connected to wireless hub using wireless trunks. Each trunk access point in the base station provides a point to multi-point wireless radio link to the remote micro access points (R-AP in figure). The remote access points provide air link service to end systems. The wireless hub and the trunk access points are co-located in the base station and connected together via 802.3 links. This figure also shows remote access points 82R connected to the wireless hub via point to point T1 links. In this scenario, no trunk APs are required.

12

To support all of the above cell architectures and the different types of access points that each cell might use, the network architecture follows the following rules:

1. Access points function as MAC layer bridges. Remote access points perform MAC bridging between the air link to the end systems and the wireless or T1 trunk to the cell site. Local access points perform MAC bridging between the air link to the end systems and the wireless hub.
2. Trunk access points also function as MAC layer bridges. They perform MAC bridging between the trunk (which goes to the access points) and the wireless hub.
3. The wireless hub is connected to all co-located MAC bridges (i.e. local access points or trunk access points) using a 802.3 link initially.

Additionally, where local access points or remote access points with T1 trunks are used, the following rules are followed.

1. Local access points are co-located with the wireless hub and connected to it using point to point 802.3 links or a shared 802.3 network. Remote access points are connected to the wireless hub using point to point T1 trunks.
2. Sectorization is supported by adding access points with sectorized antennas to the cell site.
3. For each access point connected to the wireless hub, there is a foreign agent executing in the wireless hub which participates in end system registration. MAC layer association procedures are used to keep the MAC address filter tables of the access points up to date and to perform MAC layer bridging efficiently. The wireless hub participates in MAC association functions so that only valid MAC addresses are added to the MAC address filter tables of the access points.
4. The foreign agent in the wireless hub relays frames from the access points to the MSC IWF and vice versa using the xtunnel protocol. The MAC address filter table is used to filter out those unicast MAC data frames whose MAC addresses are not present in the table. The APs always forward MAC broadcast frames and MAC frames associated with end system registration functions regardless of the contents of the MAC address filter table.
5. Local access points use ARP to resolve MAC addresses for routing IP traffic to the wireless hub. Conversely, the wireless hub also uses ARP to route IP packets to access points. UDP/IP is used for network management of access points.
6. Remote access points connected via T1 do not use ARP since the link will be a point to point link.
7. Support for hand-offs is done with assistance from the MAC layer.

In a cell architecture using wireless trunks and trunk APs, the following rules are followed.

1. Trunk access points are co-located with the wireless hub and connected to it using point to point 802.3 links or other suitable means.
2. Wireless trunk sectorization is supported by adding trunk access points with sectorized antennas to the cell site.
3. Hand-offs across backhaul sectors are done using the foreign agent in the wireless hub. For each backhaul sector, there is a foreign agent executing in the wireless hub.
4. The trunk APs do not need to participate in MAC layer end system association and hand off procedures. Their MAC address filter tables will be dynamically programmed by the wireless hub as end systems register with the network. The MAC address filter table is used to filter out unicast MAC frames. Broadcast MAC frames or MAC frames containing registration packets are allowed to always pass through.

US 6,421,714 B1

13

5. Trunk APs use ARP to resolve MAC addresses for routing IP traffic to the wireless hub. Conversely, the wireless hub use ARP to route IP packets to trunk APs. UDP/IP is used for network management of trunk APs.
6. In a single wireless trunk sector, MAC association and hand-offs from one access point to another is done using the MAC layer with the assistance of the foreign agent in the wireless hub. Using these MAC layer procedures, end systems associate with access points. As end systems move from one access point to another access point, the access points will use a MAC hand off protocol to update their MAC address filter tables. The wireless hub at the cell site provides assistance to access points to perform this function. This assistance includes relaying MAC layer hand off messages (since access points will not be able to communicate directly over the MAC layer with each other) and authenticating the end system for MAC layer registration and hand off and for updating the MAC address filter tables of the access points.
7. The foreign agent for a wireless trunk sector is responsible for relaying frames from its trunk AP to the MSC and vice versa using the xtunnel protocol. Thus, the foreign agent for a trunk AP does not care about the location of the end system with respect to access points within that wireless trunk sector. In the down link direction, it just forwards frames from the tunnel to the appropriate trunk AP which uses MAC layer bridging to send the frames to all the remote access points attached in that backhaul sector. The access points consult their MAC address filter tables and either forward the MAC frames over the access network or drop the MAC frames. As described above, the MAC address filter tables are kept up to date using MAC layer association and hand off procedures. In the up link direction, MAC frames are forwarded by the access points to the backhaul bridge which forwards them to the foreign agent in the wireless hub using the 802.3 link.
8. ARP is not be used for sending or receiving IP packets to the remote access points. The access points determines the MAC address of the wireless hub using BOOTP procedures. Conversely, the wireless hub is configured with the MAC address of remote access points. UDP/IP is used for network management of access points and for end system association and hand off messages.

IEEE Standard 802.3 links in the cell site may be replaced by other speed links.

FIG. 7 shows the protocol stack for a local access point. At the base of the stack is physical layer PHY. Physical layer PHY carries data to and from an end system over the air using radio waves as an example. When received from an end system, the AP receives data from the physical layer and unpacks it from the MAC frames (the MAC layer). The end system data frames are then repacked into an Ethernet physical layer format (IEEE 802.3 format) where it is sent via the Ethernet link to the wireless hub. When the AP's processor receives data from the wireless hub via its Ethernet link (i.e., the physical layer), the data to be transmitted to an end system, the AP packs the data in a medium access control (MAC) format, and sends the MAC layer data to its modulator to be transmitted to the end system using the PHY layer.

In FIG. 8, the MAC and PHY layers to/from the end system of FIG. 7 are replaced by a MAC and PHY for the trunk to the cell site for a remote access point.

Specifically, for a T1 trunk, the high level data link control protocol (HDLC protocol) is preferably used over the T1.

FIG. 9 depicts the protocol stack for the wireless hub that bridges the backhaul line and the trunk to the remote access

14

point. The trunk to the remote APs are only required to support remote access points (as distinct from Ethernet coupled access points). The MAC and PHY layers for the wireless trunk to the remote APs provide a point to multi-point link so that one trunk may be used to communicate with many remote APs in the same sector.

The wireless hub bridges the trunk to the remote APs and the backhaul line (e.g., T1 or T3) to the network's mobile switching center (MSC). The protocol stack in the wireless hub implements MAC and PHY layers to the MSC on top of which is implemented an IP (Internet Protocol) layer on top of which is implemented a UDP layer (Universal Datagram Protocol, in combination referred to as UDP/IP) for network management on top of which is implemented an XTunnel protocol. The XTunnel protocol is a new format that includes aspects of mobility (e.g. as in mobile IP) and aspects of the Level 2 Tunnel Protocol (L2TP). The XTunnel protocol is used to communicate from the wireless hub to the MSC and between inter-working functions (IWFs) in different networks or the same network.

In FIG. 10, the protocol stack for the relay function in the base station for supporting remote access points is shown. The relay function includes an interface to the backhaul line (depicted as the wireless hub) and an interface to the remote AP (depicted as a trunk AP). From the point of view of the wireless hub, the trunk AP (depicted in FIGS. 7 and 10) actually behaves like the AP depicted in FIG. 7. Preferably, the base station protocol stacks are split up into a wireless hub and a trunk AP with an Ethernet in between. In an N-sector wireless trunk, there are N wireless trunk APs in the cell site and one wireless hub.

In FIG. 11, the base station protocol stack for a cell architecture using a local AP is shown. The relay function includes an interface to the backhaul line (depicted as the wireless hub) and an air link interface to the end system (depicted as an AP). From the point of view of the wireless hub, the AP (depicted in FIGS. 8 and 11) actually behaves like the trunk AP depicted in FIG. 8. Preferably, the base station protocol stacks are split up into a wireless hub and a trunk AP with an Ethernet in between. In a N-sector cell, there are N access points and a single wireless hub.

The backhaul network from the base station to the MSC has the following attributes.

1. The network is capable of routing IP datagrams between the base station and the MSC.
2. The network is secure. It is not a public internet. Traffic from trusted nodes only are allowed onto the network since the network will be used for not only transporting end system traffic, but also for transporting authentication, accounting, registration and management traffic.
3. The network has the necessary performance characteristics.

In typical application, the service provider is responsible for installing and maintaining the backhaul network on which the equipment is installed.

The base stations supports the following backhaul interfaces for communicating with the MSC.

1. Base stations support IP over PPP with HDLC links using point to point T1 or fractional T3 links.
2. Base stations support IP over frame relay using T1 or fractional T3 links.
3. Base stations support IP over AAL5/ATM using T1 or fractional T3 links.
4. Base stations support IP over Ethernet links.

Since all of the above interfaces are based on IETF standard encapsulations, commercial routers may be used in the MSC to terminate the physical links of the backhaul

US 6,421,714 B1

15

network. Higher layers are passed on and processed by the various servers and other processors.

End system registration procedures above the MAC layer are supported. In the following, end system registration procedures at the MAC layer are ignored except where they impact the layers above.

End systems may register for service on their home network or from a foreign network. In both scenarios, the end system uses a foreign agent (FA) in the base station to discover a point of attachment to the network and to register. In the former case, the FA is in the end system's home network. In the latter case, the FA is in a foreign network. In either case, the network uses an IWF in the end system's home network as an anchor point (i.e., unchanging throughout the session in spite of mobility). PPP frames to and from the end system travel via the FA in the base station to the IWF in the home network. If the end system is at home, the home IWF is directly connected by means of the xtunnel protocol to the base station. Note that the home IWF may be combined with the base station in the same node. If the end system is roaming, a serving IWF in the foreign network is connected to the home IWF over an I-interface. The serving IWF relays frames between the base station and the home IWF. Note that the home IWF may be combined with the base station in the same node. From the home IWF, data is sent to a PPP server which may reside in the same IWF or to a separate server using the L2TP protocol. The separate server may be owned and operated by a private network operator (e.g. ISP or corporate intranet) who is different from the wireless service provider. For the duration of the session, the location of the home IWF and the PPP server remains fixed. If the end system moves while connected, it will have to re-register with a new foreign agent. However, the same home IWF and PPP server continues to be used. A new xtunnel is created between the new FA and the IWF and the old xtunnel between the old foreign agent and the IWF is destroyed.

FIG. 12 shows this network configuration for two end systems A and B, both of whose home wireless network is wireless service provider A (WSP-A). One end system is registered from the home wireless network and the other from a foreign wireless network. The home IWF in WSP-A serves as the anchor point for both end systems. For both end systems, data is relayed to the home IWF. The home IWF connects to an internet service provider's PPP server owned by ISP-A. Here it is assumed that both end systems have subscribed to the same ISP. If that were not the case, then the home IWF would be shown also connected to another ISP.

Within a wireless service providers network, data between base stations and the IWF is carried using the xtunnel protocol. Data between the IWF and the PPP server is carried using Level 2 Tunneling Protocol (L2TP). Data between the serving IWF and the home IWF is carried using the I-xtunnel protocol.

In a simple scenario, for a user in their home network requiring fixed service, the home IWF function may be dynamically activated in the base station. Also, the serving IWF function may be activated for a roaming user in the base station.

Always using an IWF in the home network has its advantages and disadvantages. An obvious advantage is simplicity. A disadvantage is that of always having to relay data to and from a possibly remote home IWF. The alternative is to send all the necessary information to the serving IWF so that it may connect to the end system's ISP/intranet and for the serving IWF to send accounting information in near real time back to the accounting server in the home

16

network. This functionality is more complex to implement, but more efficient because it reduces the need to relay data over potentially long distances from the foreign network to the home network.

For example, consider a case of a user who roams from Chicago to Hong Kong. If the user's home network is in Chicago and the user registers using a wireless service provider in Hong Kong, then in the first configuration, the anchor point will be the home IWF in Chicago and all data will have to be relayed from Hong Kong to Chicago and vice versa. The home IWF in Chicago will connect to the user's ISP in Chicago. With the second configuration, the end system user will be assigned an ISP in Hong Kong. Thus, data will not always have to be relayed back and forth between Chicago and Hong Kong. In the second configuration, the serving IWF will serve as the anchor and never change for the duration of the session even if the end system moves. However, the location of the FA may change as a result of end system movement in Hong Kong.

FIG. 13 shows the second network configuration. In this figure, the home network for end system A and B is WSP-A. End system A registers from its home network, using its home IWF as an anchor point, and also connects to its ISP-A using the ISP's PPP server. End system B registers from the foreign network of WSP-B and uses a serving IWF which serves as the anchor point and connects the end system to an ISP using the ISP's PPP server. In this configuration, data for end system B does not have to be relayed from the foreign network to the home network and vice versa.

In order for this configuration to work, not only must there be roaming agreements between the home and the foreign wireless service providers, but there also must be agreements between the foreign wireless service provider and the end system's internet service provider directly or through an intermediary. In the example above, not only must the wireless service provider in Hong Kong have a business agreement with the wireless service provider in Chicago, but the WSP in Hong Kong must have a business agreement with the user's Chicago ISP and access to the Chicago ISP's PPP server in Hong Kong or a business agreement with another ISP locally in Hong Kong who has a business agreement for roaming with the user's Chicago ISP. Additionally, the WSP in Hong Kong must be able to discover these roaming relationships dynamically in order to do user authentication and accounting and to set up the appropriate tunnels.

It is difficult for those companies who are in the Internet infrastructure business to work out suitable standards in the IETF for all of these scenarios. Thus, a preferable embodiment for the present systems to implement the simpler, potentially less efficient configuration, where the IWF in the home network is always used as the anchor point. However, in the presence of suitable industry standardization of protocols for Internet roaming, the second configuration should be regarded as equivalent or alternative embodiment.

An end system will have to register with the wireless network before it can start PPP and send and receive data. The end system first goes through the FA discovery and registration phases. These phases authenticate and register the end system to the wireless service provider. Once these phases are over, the end system starts PPP. This includes the PPP link establishment phase, the PPP authentication phase and the PPP network control protocol phase. Once these phases are over, the end system is able to send and receive IP packets using PPP.

The following discussion assumes that the end system is roaming and registering from a foreign network. During the

US 6,421,714 B1

17

FA discovery phase, the end system (through its user registration agent) waits for or solicits an advertisement from the foreign agent. The user registration agent uses advertisement messages sent by a near by foreign agent to discover the identity of the FA and to register. During this phase, the user registration agent of the end system selects a FA and issues a registration request to it. The FA acting as a proxy registration agent forwards the registration request to its registration server (the registration server in the foreign WSP). The registration server uses User-Name from the user registration agent's request to determine the end system's home network, and forwards the registration request for authentication to a registration server in the home network. Upon receiving the registration request relayed by the foreign registration server, the home registration server authenticates the identity of the foreign registration server and also authenticates the identity of the end system. If authentication and registration succeeds, the home registration server selects an IWF in the home network to create an I-tunnel link between the home IWF and the serving IWF (in the foreign WSP). The IWF in the home network serves as the anchor point for the duration of the PPP session.

Once the authentication and registration phases are over, the various PPP phases will be started. At the start of PPP, an L2TP connection is created between the home IWF and requested ISP/intranet PPP server. In the PPP authentication phase, PPP passwords using Password Authentication Protocol (PAP) or Challenge Authentication Protocol CHAP are exchanged and the ISP or intranet PPP server independently authenticates the identity of the end system.

Once this succeeds, the PPP network control phase is started. In this phase, an IP address is negotiated and assigned to the end system by the PPP server and the use of TCP/IP header compression is also negotiated. When this is complete, the end system is able to send and receive IP packets using PPP to its ISP or a corporate intranet.

Note that two levels of authentication are performed. The first authentication authenticates the identity of the end system to the registration server in the home network and the identities of the foreign network and the home network to each other. To perform this function, the foreign agent forwards the end system's registration request using, for example, an IETF Radius protocol to a registration server in its local MSC in a Radius Access-Request packet. Using the end system's domain name, the foreign registration server determines the identity of the end system's home network and home registration server, and acting as a Radius proxy, encapsulates and forwards the request to the end system's home registration server. If the foreign registration server cannot determine the identity of the end system's home, it may optionally forward the Radius request to a registration server that acts like a broker (e.g. one that is owned by a consortium of wireless service providers), which can in turn proxy the Radius Access-Request to the final home registration server. If the local registration server is unable to service the registration request locally or by proxying, then it rejects the foreign agent's registration request and the foreign agent rejects the end system's registration request. Upon receiving the Radius Access-Request, the home registration server performs the necessary authentication of the identities of the foreign network and the end system. If authentication and registration succeeds, the home registration server responds with a Radius Access-Response packet to the foreign registration server which sends a response to the foreign agent so that a round trip can be completed. The registration request is rejected if the home registration server is unable to comply for any reason.

18

The second level of authentication verifies the identity of the end system to the intranet or ISP PPP server. PPP authentication, separate from mobility authentication allows the infrastructure equipment to be deployed and owned separately from the ISP.

FIG. 14 is a ladder diagram showing the registration sequence for a roaming end system. It is assumed that the PPP server and the home IWF are in the same server and L2TP is not required. Note the interactions with accounting servers to start accounting on behalf of the registering end system and also directory servers to determine the identity of the home registration server and to authenticate the end system's identity. More information on accounting, billing, roaming (between service providers) and settlement will be provided below.

MAC layer messages from the user registration agent of the end system may be used to initiate Agent Solicitation. The MAC layer messages are not shown for clarity.

In FIG. 14, the end system (mobile) initially solicits an advertisement and the foreign agent replies with an advertisement that provides the end system with information about the network to which the foreign agent belongs including a care-of-address of the foreign agent. Alternatively, this phase may be removed and all network advertisements may be done by a continuously emitted MAC layer beacon message. In this case, the network is assumed to be a foreign wireless service provider. Then, a user registration agent (in the end system) incorporates the information about the foreign agent (including the user name and other security credentials) and its network into a request and sends the request to the foreign agent. The foreign agent, as a proxy registration agent, relays the request to the foreign registration server (i.e., the registration server for the foreign wireless service provider. Then, the foreign registration server, recognizing that it is not the home directory, accesses the foreign directory server with the FDD in the foreign wireless service provider to learn how to direct the registration request to the home registration server of the wireless service provider to which the end system belongs. The foreign registration server responds with the necessary forwarding information. Then, the foreign registration server encapsulates the end system's registration request in a Radius access request and relays the encapsulated request to the home registration server of the wireless service provider to which the end system belongs. The home registration server accesses the home directory server with the HDD of the home registration server to learn at least authentication information about the foreign service provider. Optionally, the home registration server accesses the subscriber's directory to learn detail subscriber service profile information (e.g., quality of service options subscribed to, etc.). When all parties are authenticated, the home registration server sends a start IWF request to the home IWF and PPP server. The home IWF and PPP server starts the home accounting server and then sends a start IWF response to the home registration server. The home registration server then sends a Radius access response to the foreign registration server. The foreign registration server then sends a start IWF request to the serving IWF server. The serving IWF server starts the serving accounting server and then sends a start IWF response to the foreign registration server. The foreign registration server sends a registration reply to the foreign agent, and the foreign agent relays the registration reply to the end system.

A link control protocol (LCP) configuration request is sent by the end system through the foreign registration server to the home IWF and PPP server. The home IWF and

US 6,421,714 B1

19

PPP server sends an LCP configuration acknowledgment through the foreign registration server to the end system.

Similarly, a password authentication protocol (PAP) authentication request is sent to and acknowledged by the home IWF and PPP server. Alternatively, a challenge authentication protocol (CHAP) may be used to authenticate. Both protocols may be used to authenticate or this phase may be skipped.

Similarly, an IP configuration protocol (IPCP) configure request is sent to and acknowledged by the home IWF and PPP server.

The connection to the end system may be terminated because of any one of the following reasons.

1. User initiated termination. Under this scenario, the end system first terminates the PPP gracefully. This includes terminating the PPP network control protocol (IPCP) followed by terminating the PPP link protocol. Once this is done, the end system de-registers from the network followed by termination of the radio link to the access point.
2. Loss of wireless link. This scenario is detected by the modem and reported to the modem driver in the end system. The upper layers of the software are notified to terminate the stacks and notify the user.
3. Loss of connection to the foreign agent. This scenario is detected by the mobility driver in the end system. After trying to re-establish contact with a (potentially new) foreign agent and failing, the driver sends an appropriate notification up the protocol stack and also signals the modem hardware below to terminate the wireless link.
4. Loss of connection to the IWF. This is substantially the same as for loss of connection to the foreign agent.
5. Termination of PPP by IWF or PPP server. This scenario is detected by the PPP software in the end system. The end system's PPP driver is notified of this event. It initiates de-registration from the network followed by termination of the wireless link to the access point.

End system service configuration refers to the concept of configuring the network service for an end system based on the subscriber's service profile. The subscriber's service profile is stored in a subscriber directory. The service profile contains information to enable the software to customize wireless data service on behalf of the subscriber. This includes information to authenticate the end system, allow the end system to roam and set up connections to the end system's internet service provider. Preferably, this information also includes other parameters, like, quality of service. In addition to the subscriber directory, a home domain directory (HDD) and a foreign domain directory (FDD) are used for roaming and for authenticating the foreign and home registration servers to each other. The HDD stores information about the end system's home network and the FDD stores information about foreign networks that a subscriber may visit.

FIG. 15 shows how these directories map into the network architecture and are used during registration for an end system that is registering at home. In step 0 the end system (mobile) solicits and receives an advertisement from the foreign agent to provides the end system with information about the network to which the foreign agent belongs. In this case, the network is the home wireless service provider. In step 1, user registration agent (in the end system) incorporates the information about the foreign agent and its network and its security credentials into a request and sends the request to the foreign agent. In step 2, the foreign agent, as a proxy registration agent, relays the request to the home registration server. In step 3, the home registration server

20

accesses the HDD of the home wireless service provider to learn at least authentication information. In step 4, the home registration server accesses the subscriber directory to learn detail subscriber service profile information (e.g., quality of service options subscribed to, etc.). In step 5, the home registration server notifies the foreign agent of the access response. In steps 6 and 7, the foreign agent notifies the end system (i.e., mobile) of the registration reply.

FIG. 16 shows directory usage for an end system that is registering from a foreign network. In step 0 the end system (mobile) solicits and receives an advertisement and the foreign agent advertises which provides the end system with information about the network to which the foreign agent belongs. In this case, the network is a foreign wireless service provider. In step 1, user registration agent (in the end system) incorporates the information about the foreign agent and its network and its security credential into a request and sends the request to the foreign agent. In step 2, the foreign agent, as a proxy registration agent, relays the request to the foreign registration server (i.e., the registration server for the foreign wireless service provider. In step 3, the foreign registration server accesses the HDD of foreign wireless service provider to learn the network to which the end system belongs. In step 4, the foreign registration server forwards the end system's request to the home registration server of the end system's home wireless service provider. In step 5, the home registration server accesses the FDD of the home registration server to learn at least authentication information about the foreign service provider. In step 6, the home registration server accesses the subscriber's directory to learn detail subscriber service profile information (e.g., quality of service options subscribed to, etc.). In step 7, the home registration server notifies the foreign registration server of the access response. In step 8, the foreign registration server forwards to the foreign agent the access response. In step 9, the foreign agent notifies the end system (i.e., mobile) of the registration reply.

Protocol handling scenarios handle bearer data and the associated stacks for transporting bearer data to and from an end system. The protocol stacks for the cell architectures use local APs (FIG. 17) and remote APs (FIG. 18).

FIG. 17 shows the protocol stacks for handling communications between an end system (in its home network) and a home IWF for End System@Home. FIG. 17 shows the protocol handling for a cell architecture where the access point and the wireless hub are co-located.

FIG. 18 shows the protocol handling for a cell architecture where the access point is located remotely from the wireless hub. As shown, PPP terminates in the IWF and the configuration provides direct internet access. The configuration for the case where the PPP server is separate from the IWF is described later.

In FIG. 18, PPP frames from the end system are encapsulated in RLP (radio link protocol) frames which are encapsulated at the remote access point in MAC frames for communicating with the trunk access point (i.e., an access point physically located near the wireless hub), the remote access point being coupled to the access point by, for example, a wireless trunk). The access point functions as a MAC layer bridge and relays frames from the air link to the foreign agent in the wireless hub. The foreign agent de-encapsulates the RLP frames out of the MAC frames, and using the xtunnel protocol, relays the RLP frames to the IWF. A similar, albeit reverse, process occurs for transmitting frames from the IWF to the end system.

If the end system moves to another foreign agent, then a new xtunnel will be automatically created between the new

US 6,421,714 B1

21

foreign agent and the IWF, so that PPP traffic continues to flow between them, without interruption.

In the remote AP cell architecture (FIG. 18) using wireless trunks between the remote AP and the trunk AP, the air link between the end system and the access point may operate at a different frequency (f1) and use a different radio technology as compared to the frequency (f2) and radio technology of the trunk.

FIG. 19 shows the protocol stacks for a roaming end system. The serving IWF uses of the I-xtunnel protocol between the serving IWF and home IWF. The rest of the protocol stacks remain unchanged and are not shown. This architecture may be simplified by merging the serving IWF into the base station, thus eliminating the XWD protocol.

The RLP layer uses sequence numbers to drop duplicate PPP datagrams and provide in-sequence delivery of PPP datagrams between the end system and the IWF. It also provides a configurable keep-alive mechanism to monitor link connectivity between the end system and the IWF. Additionally, in an alternative embodiment, the RLP layer also provides re-transmission and flow control services in order to reduce the overall bit error rate of the link between the end system and the IWF. The RLP between the end system and the IWF is started at the beginning of the session and remains active throughout the session and even across hand-offs.

In contrast to the specification in the mobile IP RFC (RFC 2003), IP in IP encapsulation is not used for tunneling between the foreign agent and the home IWF. Instead a new tunneling protocol, implemented on top of UDP is used. This tunneling protocol is a simplified version of the L2TP protocol. The reasons for this choice are as follows.

1. The encapsulation protocol specified in RFC 2003 does not provide flow control or in-sequence delivery of packets. The presently described network may need these services in the tunnel over the backhaul. Flow control may be needed to reduce the amount of re-transmissions over the air link because of packet loss due to flow control problems over the network between the base station and the MSC or because of flow control problems in the base station or the IWF.
2. By using a UDP based tunneling protocol, the implementation can be done at the user level and then put into the kernel for performance reasons, after it has been debugged.
3. Using RFC 2003, there is no easy way of creating tunnels taking into account quality of service and load balancing. In order to take QOS into account, it should be possible to set up tunnels over links that already provide the required QOS. Secondly, using RFC 2003, there is no easy way to provide load balancing to distribute bearer traffic load over multiple links between the base station and the MSC.
4. In order to implement IP in IP encapsulation as specified in RFC 2003, developers require access to IP source code. In commercial operating systems, source code for the TCP/IP stack is generally proprietary to other equipment manufacturers. Purchasing the TCP/IP stack from a vendor and making changes to the IP layer to support mobile IP tunneling would require a developer to continue supporting a variant version of the TCP/IP stack. This adds cost and risk.

While it is noted that the tunneling protocol between the base station and the IWF is non-standard and that the wireless service provider will not be able to mix and match equipment from different vendors, the use of a non-standard tunneling protocol within a single wireless service provider network is transparent to end systems and equipment from other vendors.

22

The new tunneling protocol is based on L2TP. By itself, L2TP is a heavyweight tunneling protocol so that L2TP has a lot of overhead associated with tunnel creation and authentication. The new tunneling protocol of the present system has less overhead. The new xtunnel protocol has the following features.

1. The xtunnel creation adds vendor specific extensions to Radius Access Request and Radius Access Response messages between the base station and the registration server. These extensions negotiate tunnel parameters and to create the tunnel.
2. The registration server is able to delegate the actual work of tunneling and relaying packets to a different IP address, and therefore, to a different server in the MSC. This permits the registration server to do load balancing across multiple IWF servers and to provide different QOS to various users.
3. The xtunnel protocol supports in-band control messages for tunnel management. These messages include echo request/response to test tunnel connectivity, disconnect request/response/notify to disconnect the tunnel and error notify for error notifications. These messages are sent over the tunneling media, for example, UDP/IP.
4. The xtunnel protocol sends payload data over the tunneling media, for example, UDP/IP. The xtunnel protocol supports flow control and in-sequence packet delivery.
5. The xtunnel protocol may be implemented over media other than UDP/IP for quality of service.

The network supports direct internet connectivity by terminating the PPP in the home IWF and routing IP packets from the IWF to the internet via a router using standard IP routing techniques. Preferably, the IWF runs Routing Information Process (RIP), and the router also runs RIP and possibly other routing protocols like Open Shortest Path First (OSPF).

The network supports a first configuration for a wireless service provider who is also an internet service provider. In this configuration, the home IWF in the MSC also functions as a PPP server. This IWF also runs internet routing protocols like RIP and uses a router to connect to the internet service provider's backbone network.

The network supports a second configuration for a wireless service provider who wishes to allow end systems to connect to one or more internet service providers, either because the WSP itself is not ISPs, or because the WSP has agreements with other ISPs to provide access to end users. For example, a wireless service provider may elect to offer network access to an end user and may have an agreement with a 3rd party ISP to allow the user who also has an account with the 3rd party ISP to access the ISP from the WSP network. In this configuration, the PPP server does not run in the home IWF installed at the MSC. Instead, a tunneling protocol like L2TP (Layer Two Tunneling Protocol) is used to tunnel back to the ISP's PPP server. FIG. 10 shows the protocol stacks for this configuration for an end system that is at home.

The location of the home IWF and the ISP PPP server remains fixed throughout the PPP session. Also, the L2TP tunnel between the IWF and the ISP's PPP server remains up throughout the PPP session. The physical link between the IWF and the PPP server is via a router using a dedicated T1 or T3 or frame relay or ATM network. The actual nature of the physical link is not important from the point of view of the architecture.

This configuration also supports intranet access. For intranet access, the PPP server resides in the corporate intranet and the home IWF uses L2TP to tunnel to it.

US 6,421,714 B1

23

For a fixed end system, the protocol handling for intranet or ISP access is as shown in FIG. 20 with the difference that the roaming end system uses a serving IWF to connect to its home IWF. The protocol handling between a serving IWF and a home IWF has been described earlier. In FIG. 20, the home IWF may be merged into the wireless hub eliminating the X-Tunnel protocol. Also, the serving IWF may be merged into the wireless hub, thus eliminating the X-Tunnel protocol.

FIG. 21 shows the protocol stacks used during the registration phase (end system registration) for a local AP cell architecture. The stack for a remote AP cell architecture is very similar.

The scenario shown above is for a roaming end system. For an end system at home, there is no foreign registration server in the registration path.

Note the mobility agent in the end system. The mobility agent in the end system and foreign agent in the wireless hub are conceptually similar to the mobile IP RFC 2002. The mobility agent handles network errors using time-outs and re-tries. Unlike the known protocol stacks for bearer data, RLP is not used. The foreign agent and the registration servers use Radius over UDP/IP to communicate with each other for registering the end system.

Several aspects of security must be considered. The first, authenticating the identities of the end system and the foreign/home networks during the wireless registration phase. Second, authenticating the identity of the end system with its PPP server during the PPP authentication phase. Third, authentication for storing accounting data, for billing and for updating home domain information. Fourth, encryption of bearer traffic transmitted to and from the end system. Fifth, encryption for exchanging billing information across service provider boundaries.

Shared secrets are used to authenticate the identity of end systems with their home networks and the identity of the home and foreign networks with each other during wireless registration.

End system authentication uses a 128-bit shared secret to create an authenticator for its registration request. The authenticator is created using the known MD5 message digest algorithm as described in the mobile IP RFC 2002. Alternatively, a different algorithm may be used. The shared secret is not sent in the registration request by the end system. Only the authenticator is sent. On receiving the registration request from the end system, the home registration server re-computes the authenticator over the registration request data using the shared secret. If the computed authenticator value matches the authenticator value sent by the end system, the home registration server allows the registration process to proceed. If the values do not match, the home registration server logs the event, generates a security violation alarm and a nak (i.e., a negative acknowledgment) to the request.

In the registration reply, the home registration server does the same—that is to say, uses the shared secret to create an authenticator for the registration reply that it sends to the end system. Upon receiving the reply, the end system re-computes the authenticator using the shared secret. If the computed value does not match the authenticator value sent by the home registration server in the reply, the end system discards the reply and tries again.

These network security concepts are similar to the concepts defined in mobile IP RFC 2002. According to the RFC, a mobility security association exist between each end system and its home network. Each mobility security association defines a collection of security contexts. Each secu-

24

urity context defines an authentication algorithm, a mode, a secret (shared or public-private), style of replay protection and the type of encryption to use. In the context of the present network, the end system's User-Name (in lieu of the mobile IP home address) is used to identify the mobility security association between the end system and its home network. Another parameter, called the security parameter index (SPI), is used to select a security context within the mobility security association. In a basic embodiment of the invention, only the default mobile IP authentication algorithm (keyed-MD5) and the default mode ("prefix+suffix") are supported with 128-bit shared secrets. Network users are allowed to define multiple shared secrets with their home networks. The mechanism for creating security contexts for end users, assigning an SPI to each security context and for setting the contents of the security context (which includes the shared secret) and for modifying their contents are described below. During registration, a 128-bit message digest is computed by the end system in prefix+suffix mode using the MD5 algorithm. The shared secret is used as the prefix and the suffix for the data to be protected in the registration request. The authenticator thus computed, along with the SPI and the User-Name are transmitted in the registration request by the end system. Upon receiving the end system's registration request, the foreign registration server relays the request along with the authenticator and the SPI, unchanged to the home registration server. Upon receiving the registration request directly from the end system or indirectly via a foreign registration server, the home registration server uses the SPI and the User-Name to select the security context. The home server re-computes the authenticator using the shared secret. If the computed authenticator value matches the value of the authenticator sent in the request by the end system, the user's identity will have been successfully authenticated. Otherwise, the home registration server naks (negatively acknowledges) the registration request sent by the end system.

The registration reply sent by the home registration server to the end system is also authenticated using the algorithm described above. The SPI and the computed authenticator value is transmitted in the registration reply message by the home server to the end system. Upon receiving the reply, the end system re-computes the authenticator, and if the computed value does not match the transmitted value, it will discard the reply and retry.

The user's end system has to be configured with the shared secret and SPIs for all security contexts that the user shares with its registration server(s). This configuration information is preferably stored in a Win 95 registry for Windows 95 based end systems. During registration, this information is accessed and used for authentication purposes.

In the network, Radius protocols are used by foreign agent FA to register the end system and to configure the xtunnel between the wireless hub and the home and serving IWFs on behalf of the end system. On receiving a registration request from the end system, the FA creates a Radius Access-Request packet, stores its own attributes into the packet, copies the end system's registration request attributes unchanged into this packet and sends the combined request to the registration server in the MSC.

Radius authentication requires that the Radius client (in this case, the FA in the base station) and the Radius server (in this case, the registration server in the MSC) share a secret for authentication purposes. This shared secret is also used to encrypt any private information communicated between the Radius client and the Radius server. The shared

US 6,421,714 B1

25

secret is a configurable parameter. The network follows the recommendations in the Radius RFC and uses the shared secret and the MD5 algorithm for authentication and for encryption, where encryption is needed. The Radius-Access Request packet sent by the FA contains a Radius User-Name attribute (which is provided by the end system) and a Radius User-Password attribute. The value of the User-Password attribute is also a configurable value and encrypted in the way recommended by the Radius protocol. Other network specific attributes, which are non-standard attributes from the point of view of the Radius RFC standards, are encoded as vendor specific Radius attributes and sent in the Access-Request packet.

The following attributes are sent by the FA to its registration server in the Radius Access-Request packet.

1. User-Name Attribute. This is the end system's user-name as supplied by the end system in its registration request.
2. User-Password Attribute. This user password is supplied by the base station/wireless hub on behalf of the user. It is encoded as described in the Radius RFC using the secret shared between the base station and its registration server.
3. NAS-Port. This is the port on the base station.
4. NAS-IP-Address. This is the IP address of the base station.
5. Service-Type. This is framed service.
6. Framed-Protocol. This is a PPP protocol.
7. Xtunnel Protocol Parameters. These parameters are sent by the base station to specify the parameters necessary to set up the xtunnel protocol on behalf of the end system. This is a vendor-specific attribute.
8. AP-IP-Address. This is the IP address of the AP through which the user is registering. This is a vendor-specific attribute.
9. AP-MAC-Address. This is the MAC address of the AP through which the user is registering. This is a vendor-specific attribute.
10. End system's Registration Request. The registration request from the end system is copied unchanged into this vendor specific attribute.

The following attributes are sent to the FA from the registration server in the Radius Access-Response packet.

1. Service Type. This is a framed service.
2. Framed-Protocol. This is a PPP.
3. Xtunnel Protocol Parameters. These parameters are sent by the registration server to specify the parameters necessary to set up the xtunnel protocol on behalf of the end system. This is a vendor-specific attribute.
4. Home Registration Server's Registration Reply. This attribute is sent to the FA from the home registration server. The FA relays this attribute unchanged to the end system in a registration reply packet. If there is a foreign registration server in the path, this attribute is relayed by it to the PA unchanged. It is coded as a vendor-specific attribute.

To provide service to roaming end systems, the foreign network and the home network are authenticated to each other for accounting and billing purposes using the Radius protocol for authentication and configuration. This authentication is performed at the time of end system registration. As described earlier, when the registration server in the foreign network receives a registration request from an end system (encapsulated as a vendor specific attribute in a Radius-Access Request packet by the FA), it uses the end system's User-Name to determine the identity of the end system's home registration server by consulting its home domain directory HDD. The following information is stored in home domain directory IIDD and accessed by the foreign

26

registration server in order to forward the end system's registration request.

1. Home Registration Server IP Address. This is the IP address of the home registration server to forward the registration request.
2. Foreign Registration Server Machine Id. This is the machine ID of the foreign registration server in SMTP (simplified mail transfer protocol) format (e.g., machine@fqdn where machine is the name of the foreign registration server machine and fqdn is the fully qualified domain name of the foreign registration server's domain).
3. Tunneling Protocol Parameters. These are parameters for configuring the tunnel between the serving IWF and the home IWF on behalf of the end system. These include the tunneling protocol to be used between them and the parameters for configuring the tunnel.
4. Shared Secret. This is the shared secret to be used for authentication between the foreign registration server and the home registration server. This secret is used for computing the Radius User-Password attribute in the Radius packet sent by the foreign registration server to the home registration server. It is defined between the two wireless service providers.
5. User-Password. This is the user password to be used on behalf of the roaming end system. This user password is defined between the two wireless service providers. This password is encrypted using the shared secret as described in the Radius RFC.
6. Accounting Parameters. These are parameters for configuring accounting on behalf of the end system that is registering. These parameters are sent by the registration server to its IWF for configuring accounting on behalf of the end system.

Using this information, the foreign registration server creates a Radius Access-Request, adds its own registration and authentication information into the Radius Access-Request, copies the registration information sent by the end system unchanged into the Radius Access-Request and sends the combined request to the home registration server.

Upon receiving the Radius-Access Request from the foreign registration server (for a roaming end system) or directly from the FA (for an end system at home), the home registration server consults its own directory server for the shared secrets to verify the identity of the end system and the identity of the foreign registration server in a roaming scenario by re-computing authenticators.

After processing the request successfully, the home registration server creates a Radius Access-Accept response packet and sends it to the foreign registration server if the end system is roaming, or directly to the FA from which it received the Radius Access-Request. The response contains the registration reply attribute that the FA relays to the end system.

If the request can not be processed successfully, the home registration server creates a Radius Access-Reject response packet and sends it to the foreign registration server if the end system is roaming, or directly to the FA from which it received the Radius Access-Request. The response contains the registration reply attribute that the FA will relay to the end system.

In a roaming scenario, the response from the home registration server is received by the foreign registration server. It is authenticated by the foreign registration server using the shared secret. After authenticating, the foreign registration server processes the response, and in turn, it generates a Radius response packet (Accept or Reject) to send to the FA. The foreign registration server copies the

US 6,421,714 B1

27

registration reply attribute from the home registration server's Radius response packet, unchanged, into its Radius response packet.

When the FA receives the Radius Access-Response or Radius Access-Reject response packet, it creates a registration reply packet using the registration reply attributes from the Radius response, and sends the reply to the end system, thus completing the round trip registration sequence.

Mobile IP standards specifies that replay protection for registrations are implemented using time stamps, or optionally, using nonces. However, since replay protection using time stamps requires adequately synchronized time-of-day clocks between the corresponding nodes, the present system implements replay protection during registration using nonces even though replay protection using time stamps is mandatory in the Mobile IP standards and the use of nonces is optional. However, replay protection using time stamps as an alternative embodiment is envisioned.

The style of replay protection used between nodes is stored in the security context in addition to the authentication context, mode, secret and type of encryption.

The network supports the use of PPP PAP (password authentication) and CHAP (challenge authenticated password) between the end system and its PPP server. This is done independently of the registration and authentication mechanisms described earlier. This allows a private intranet or an ISP to independently verify the identity of the user.

Authentication for accounting and directory services is described below with respect to accounting security. Access to directory servers from network equipment in the same MSC need not be authenticated.

The network supports encryption of bearer data sent between the end system and the home IWF. End systems negotiate encryption to be on or off by selecting the appropriate security context. Upon receiving the registration request, the home registration server grants the end system's request for encryption based upon the security context. In addition to storing the authentication algorithm, mode, shared secret and style of replay protection, the security context is also used to specify the style of encryption algorithm to use. If encryption is negotiated between the end system and the home agent, then the complete PPP frame is so encrypted before encapsulation in RLP.

The IWF, the accounting server and the billing system are part of the same trusted domain in the MSC. These entities are either connected on the same LAN or part of a trusted intranet owned and operated by the wireless service provider. Transfer of accounting statistics between the IWF and the accounting server and between the accounting server and the customer's billing system may be encrypted using Internet IP security protocols like IP-Sec.

The network makes it more difficult to monitor the location of the end system because it appears that all PPP frames going to and from the end system go through the home IWF regardless of the actual location of the end system device.

Accounting data is collected by the serving IWF and the home IWF in the network. Accounting data collected by the serving IWF is sent to an accounting server in the serving IWF's MSC. Accounting data collected by the home IWF is sent to an accounting server in the home IWF's MSC. The accounting data collected by the serving IWF is used by the foreign wireless service provider for auditing and for settlement of bills across wireless service provider boundaries (to support roaming and mobility). The accounting data collected by the home IWF is used for billing the end user and also for settlement across wireless service provider boundaries to handle roaming and mobility.

28

Since all data traffic flows through the home IWF, regardless of the end system's location and the foreign agent's location, the home IWF has all the information to generate bills for the customer and also settlement information for the use of foreign networks.

The serving IWF and the home IWF preferably use the Radius accounting protocol for sending accounting records for registered end systems. The Radius accounting protocol is as documented in a draft IETF RFC. For the present invention, the protocol has to be extended by adding vendor specific attributes for the network and by adding check-pointing to the Radius Accounting protocol. Check-pointing in this context refers to the periodic updating of accounting data to minimize risk of loss of accounting records.

The Radius accounting protocol runs over UDP/IP and uses re-tries based on acknowledgment and time outs. The Radius accounting client (serving IWFs or home IWFs) send UDP accounting request packets to their accounting servers which send acknowledgments back to the accounting clients.

In the network, the accounting clients (serving IWF and the home IWF) emit an accounting start indication at the start of the user's session and an accounting stop indication at the end of the user's session. In the middle of the session, the accounting clients emit accounting checkpoint indications. In contrast, the Radius accounting RFC does not specify an accounting checkpoint indication. The software of the present system creates a vendor specific accounting attribute for this purpose. This accounting attribute is present in all Radius Accounting-Request packets which have Acct-Status-Type of Start (accounting start indications). The value of this attribute is used to convey to the accounting server whether the accounting record is a check-pointing record or not. Check-pointing accounting reports have a time attribute and contain cumulative accounting data from the start of the session. The frequency of transmitting check-point packets is configurable in the present invention.

The serving IWF and the home IWF are configured by their respective registration servers for connecting to their accounting servers during the registration phase. The configurable accounting parameters include the IP address and UDP port of the accounting server, the frequency of check-pointing, the session/multi-session id and the shared secret to be used between the accounting client and the accounting server.

The network records the following accounting attributes for each registered end system. These accounting attributes are reported in Radius accounting packets at the start of the session, at the end of the session and in the middle (check-point) by accounting clients to their accounting servers.

1. User Name. This is like the Radius User-Name attribute discussed above. This attribute is used to identify the user and is present in all accounting reports. The format is "user@domain" where domain is the fully qualified domain name of the user's home.
2. NAS IP Address. This is like the Radius NAS-IP-Address attribute discussed above. This attribute is used to identify the IP address of the machine running the home IWF or the serving IWF.
3. Radio Port. This attribute identifies the radio port on the access point providing service to the user. This attribute is encoded as a vendor specific attribute.
4. Access Point IP Address. This attribute identifies the IP address of the access point providing service to the user. This attribute is encoded as a vendor specific attribute.
5. Service Type. This is like the Radius Service-Type attribute described above. The value of this attribute is Framed.

US 6,421,714 B1

29

6. Framed Protocol. This is like the Radius Framed-Protocol attribute described above. The value of this attribute is set to indicate PPP.
7. Accounting Status Type. This is like the Radius Acct-Status-Type attribute described above. The value of this attribute may be Start to mark the start of a user's session with the Radius client and Stop to mark the end of the user's session with the Radius client. For accounting clients, the Acct-Status-Type/Start attribute is generated when the end system registers. The Acct-Status-Type/Stop attribute is generated when the end system de-registers for any reason. For checkpoints, the value of this attribute is also Start and the Accounting Checkpoint attribute is also present.
8. Accounting Session Id. This is like the Radius Acct-Session-Id described above. In a roaming scenario, this session id is assigned by the foreign registration server when the end system issues a registration request. It is communicated to the home registration server by the foreign registration server during the registration sequence. The home network and the foreign network both know the Acct-Session-Id attribute and are able to emit this attribute while sending accounting records to their respective accounting servers. In a "end system-at-home" scenario, this attribute is generated by the home registration server. The registration server communicates the value of this attribute to the IWF which emits it in all accounting records.
9. Accounting Multi-Session Id. This is like the Radius Acct-Multi-Session-Id discussed above. This id is assigned by the home registration server when a registration request is received from a FA directly or via a foreign registration server on behalf of an end system. It is communicated to the foreign registration server by the home registration server in the registration reply message. The registration server(s) communicates the value of this attribute to the IWF(s) which emit it in all accounting records.

With true mobility added to the architecture, the id is used to relate together the accounting records from different IWFs for the same end system if the end system moves from one IWF to another. For hand-offs across IWF boundaries, the Acct-Session-Id is different for accounting records emanating from different IWFs. However, the Acct-Multi-Session-Id attribute is the same for accounting records emitted by all IWFs that have provided service to the user. Since the session id and the multi-session id are known to both the foreign network and the home network, they are able to emit these attributes in accounting reports to their respective accounting servers. With the session id and the multi-session id, billing systems are able to correlate accounting records across IWF boundaries in the same wireless service provider and even across wireless service provider boundaries.

1. Accounting Delay Time. See Radius Acct-Delay-Time attribute.
2. Accounting Input Octets. See Radius Acct-Input-Octets. This attribute is used to keep track of the number of octets sent by the end system (input to the network from the end system). This count is used to track the PPP frames only. The air link overhead, or any overhead imposed by RLP, etc. is not counted.
3. Accounting Output Octets. See Radius Acct-Output-Octets. This attribute is used to keep track of the number of octets sent to the end system (output from the network to the end system). This count is used to track the PPP frames only. The air link overhead, or any overhead imposed by RLP, etc. and is not counted.

30

4. Accounting Authentic. See Radius Acct-Authentic attribute. The value of this attribute is Local or Remote depending on whether the serving IWF or the home IWF generates the accounting record.
5. Accounting Session Time. See Radius Acct-Session-Time attribute. This attribute indicates the amount of time that the user has been receiving service. If sent by the serving IWF, this attribute tracks the amount of time that the user has been receiving service from that serving IWF. If sent by the home IWF, this attribute tracks the amount of time that the user has been receiving service from the home IWF.
6. Accounting Input Packets. See Radius Acct-Input-Packets attribute. This attribute indicates the number of packets received from the end system. For a serving IWF, this attribute tracks the number of PPP frames input into the serving IWF from an end system. For a home IWF, this attribute tracks the number of PPP frames input into the home IWF from an end system.
7. Accounting Output Packets. See Radius Acct-Output-Packets attribute. This attribute indicates the number of packets sent to the end system. For a serving IWF, this attribute tracks the number of PPP frames output by the serving IWF to the end system. For a home IWF, this attribute tracks the number of PPP frames sent to the end system from the home IWF.
8. Accounting Terminate Cause. See Radius Acct-Terminate-Cause attribute. This attribute indicates the reason why a user session was terminated. In addition, a specific cause code is also present to provide additional details. This attribute is only present in accounting reports at the end of the session.
9. Network Accounting Terminate Cause. This attribute indicates a detailed reason for terminating a session. This specific attribute is encoded as a vendor specific attribute and is only reported in a Radius Accounting attribute at the end of session. The standard Radius attribute Acct-Terminate-Cause is also present. This attribute provides specific cause codes, not covered by the Acct-Terminate-Cause attribute.
10. Network Air link Access Protocol. This attribute indicates the air link access protocol used by the end system. This attribute is encoded as a vendor specific attribute.
11. Network Backhaul Access Protocol. This attribute indicates the backhaul access protocol used by the access point to ferry data to and from the end system. This attribute is encoded as a vendor specific attribute.
12. Network Agent Machine Name. This attribute is the fully qualified domain name of the machine running the home IWF or the serving IWF. This specific attribute is encoded in vendor specific format.
13. Network Accounting Check-point. Since the Radius accounting RFC does not define a check-point packet, the present network embodiment uses a Radius accounting start packet with this attribute to mark a check-point. The absence of a check-point attribute means a conventional accounting start packet. The presence of this attribute in a accounting start packet means a accounting check-point packet. Accounting stop packets do not have this attribute. In the preferred embodiment, every accounting packet and the corresponding reply must be authenticated using MD5 and a shared secret. The IWFs are configured with a shared secret that are used by them for authentication during communication with their Radius accounting server. The shared secrets used by the IWFs for communicating with accounting servers are stored in the home/foreign domain directory located in the MSC. The shared secrets for

US 6,421,714 B1

31

accounting security are communicated to the IWFs by their registration servers during the end system registration sequence.

The accounting server software runs in a computer located in the MSC. The role of the accounting server in the system is to collect raw accounting data from the network elements (the home and the serving IWFs), process the data and store it for transfer to the wireless service provider's billing system. The accounting server does not include a billing system. Instead, it includes support for an automatic or manual accounting data transfer mechanism. Using the automatic accounting data transfer mechanism, the accounting server transfers accounting records in AMA billing format to the customer's billing system over a TCP/IP transport. For this purpose, the system defines AMA billing record formats for packet data. Using the manual transfer mechanism, customers are able to build a tape to transfer accounting records to their billing system. In order to build the tape to their specifications, customers are provided with information to access accounting records so that they may process them before writing them to tape.

In FIG. 22, the raw accounting data received by the accounting server from the home or serving IWFs are processed and stored by the accounting server. The processing done by the accounting server includes filtering, compression and correlation of the raw accounting data received from the IWF. A high availability file server using dual active/standby processors and hot swappable RAID disks is used for buffering the accounting data while it is transiting through the accounting server.

The accounting server delays processing of the raw accounting data until an end system has terminated its session. When an end system terminates its session, the accounting server processes the raw accounting data that it has collected for the session and stores an accounting summary record in a SQL database. The accounting summary record stored in the SQL data base points to an ASN.1 encoded file. This file contains detailed accounting information about the end system's session. The data stored in the accounting server is then transferred by the billing data transfer agent to the customer's billing system. Alternatively, the wireless service provider may transfer the accounting data from the SQL database and/or the ASN.1 encoded file to the billing system via a tape. The data base scheme and the format of the ASN.1 encoded file are documented and made available to customers for this purpose. If the volume of processed accounting data stored in the accounting system exceeds a high water mark, the accounting server generates an NMS alarm. This alarm is cleared when the volume of data stored in the accounting server falls below a low water mark. The high and low water marks for generating and clearing the alarm are configurable. The accounting server also generates an NMS alarm if the age of the stored accounting data exceeds a configurable threshold. Conversely, the alarm is cleared, when the age of the accounting data falls below the threshold.

The subscriber directory is used to store information about subscribers and is located in the home network. The home registration server consults this directory during the registration phase to authenticate and register an end system. For each subscriber, the subscriber directory stores the following information.

1. User-Name. This field in the subscriber record will be in SMTP format (e.g., user@fqdn) where the user sub-field will identify the subscriber in his or her wireless home domain and the fqdn sub-field will identify the wireless home domain of the subscriber. This field is sent by the

32

end system in its registration request during the registration phase. This field is assigned by the wireless service provider to the subscriber at the time of subscription to the network service. This field is different than the user name field used in PPP.

2. Mobility Security Association. This field in the subscriber record contains the mobility security association between the subscriber and his or her home network. As described above, a mobility security association exists between each subscriber and its home registration server. The mobility security association defines a collection of security contexts. Each security context defines an authentication algorithm, an authentication mode, a shared secret, style of replay protection and the type of encryption (including no encryption) to use between the end system and its home server. During registration, the home registration server retrieves information about the subscriber's security context from the subscriber directory using the User-Name and the security parameter index (SPI) supplied by the end system in its registration request. The information in the security context is used for enforcing authentication, encryption and replay protection during the session. The mobility security association is created by the wireless service provider at the time of subscription. It is up to the wireless service provider to permit the subscriber to modify this association either by calling up a customer service representative or by letting subscribers access to a secure Web site. The Web site software will export web pages which the wireless service provider may make accessible to subscribers from a secure web server. In this way, subscribers are able to view/modify the contents of the mobility security association in addition to other subscriber information that the service provider may make accessible.
3. Modem MAC Address. This field contains the MAC address of the modem owned by the subscriber. In addition to the shared secret, this field is used during registration to authenticate the user. It is possible to turn off MAC address based authentication on a per user basis. The MAC address is communicated to the home registration server during registration.
4. Enable MAC Address Authentication. This field is used to determine if MAC address based authentication is enabled or disabled. If enabled, the home registration server checks the MAC address of the registering end system against this field to validate the end system's identity. If disabled, then no checking is done.
5. Roaming Enabled Flag. If this field is set to enabled, then the end system is allowed to roam to foreign networks. If this field is disabled, then the end system is not permitted to roam to foreign networks.
6. Roaming Domain List. This field is meaningful only if the Roaming Enabled Flag is set to enabled. This field contains a list of foreign domains that the end system is allowed to roam to. When the contents of this list is null and the Roaming Enabled Flag is set to enabled, the end system is allowed to roam freely.
7. Service Enable/Disable Flag. This field may be set to disabled by the system administrator to disable service to a subscriber. If this field is disabled, then the subscriber is permitted to register for service. If the subscriber is registered and the value of this field is set to disabled, then the subscriber's end system is immediately disconnected by the network.
8. Internet Service Provider Association. This field contains information about the subscriber's internet service provider. This information is used by the IWF during the PPP

US 6,421,714 B1

33

registration phase to perform authentication with the internet service provider on behalf of the end system and also to create a L2TP tunnel between the IWF and the internet service provider's PPP server. This field contains the identity of the subscriber's ISP. The IWF uses this information to access the ISP directory for performing authentication and setting up the L2TP tunnel on behalf of the end system.

9. Subscriber's Name & Address Information. This field contains the subscriber's name, address, phone, fax, e-mail address, etc.

The home domain directory (HDD) is used by the registration server to retrieve parameters about the end system to complete registration on behalf of the end system. Using this information, the registration server determines if the end system is registering at home or if the end system is a roaming end system. In the former case, the registration server assumes the role of a home registration server and proceed with end system registration. In the latter case, the registration server assumes the role of a foreign registration server and, acting as a Radius proxy, it forwards the request to the real home registration server whose identity it gets from this directory. For roaming end system, the parameters stored in the HDD include the IP address of the home registration server, the home-foreign shared secret, the home-serving IWF tunnel configuration etc. The HDD is located in the MSC.

The following information is stored in the HDD.

1. Home Domain Name. This field is used as the key to search the HDD for an entry that matches the fully qualified home domain name provided by the end system in its registration request.
2. Proxy Registration Request. This field is used by the registration server to determine if it should act as a foreign registration server and proxy the end system's registration request to the real home registration server.
3. Home Registration Server DNS Name. If the proxy registration request flag is TRUE, this field is used to access the DNS name of the real home registration server. Otherwise, this field is ignored. The DNS name is translated to an IP address by the foreign registration server. The foreign registration server uses the IP address to relay the end system's registration request.
4. Foreign Domain Name. If the proxy registration request flag is TRUE, this field is used to identify the foreign domain name to the end system's home registration server. Otherwise, this field is ignored. The foreign registration server uses this information to create the foreign server machine id in SMTP format, for example, machine@fqdn. This machine id is sent to the home registration server by the foreign registration server in the Radius-Access Request.
5. Shared Secret. If the proxy registration request flag is TRUE, the shared secret is used between the foreign and home registration servers to authenticate their identity with each other. Otherwise this field is ignored.
6. Tunneling Protocol Parameters. This field is used to store parameters to configure the tunnels to provide service to the end system. For an end system at home, this includes information on tunnel parameters between the base station and the home IWF and from the home IWF to the PPP server. For a roaming end system, this includes tunneling parameters from the base station to the serving IWF and from the serving IWF to the home IWF. At a minimum, for each tunnel, this field contains the type of tunneling protocol to use and any tunneling protocol specific parameters. For example, this field may contain the identifier for

34

the tunneling protocol L2TP and any additional parameters required to configure the L2TP tunnel between the IWF and its peer.

7. Accounting Server Association. This field is used to store information needed by the IWF to generate accounting data on behalf of the end system. It contains the name of the accounting protocol (e.g. RADIUS), the DNS name of the accounting server and additional parameters specific to the accounting protocol like the UDP port number, the shared secret that the IWF must use in the Radius Accounting protocol, the frequency of check-pointing, the seed for creating the session/multi-session id, etc. The accounting server's DNS name is translated to the accounting server's IP address, which is sent to the IWF.

For wireless service providers that have roaming agreements with each other, the HDD is used for authentication and to complete the registration process. If an end system roams from its home network to a foreign network, the foreign registration server in that network consults the HDD in its MSC to get information about the visiting end system's home registration and to authenticate the home network before it provides service to the visiting end system.

The software for home domain directory management preferably provides a graphical user interface (GUI) based HDD management interface for system administrators. Using this GUI, system administrators are able to view and update entries in the HDD. This GUI is not intended for use by foreign wireless network service providers to perform remote updates based on roaming agreements. It is only intended for use by trusted personnel of the home wireless service provider operating behind fire walls.

The foreign domain directory (FDD) provides functionality that is the reverse of the home domain directory. The FDD is used by the home registration server to retrieve parameters about the foreign registration server and the foreign network in order to authenticate the foreign network and create a tunnel between a serving IWF and a home IWF. These parameters include the home-foreign shared secret, the home IWF-serving IWF tunnel configuration, etc. The FDD is preferably located in the home registration server's MSC. The FDD is used by home registration servers for registering roaming end systems.

The following information will be stored in the FDD.

1. Foreign Domain Name. This field is used as the key to search the FDD for an entry that matches the fully qualified domain name of the foreign registration server relaying the registration request.
2. Shared Secret. This is the shared secret used between the foreign and home registration servers to authenticate their identity mutually with each other.
3. Home IWF-Serving IWF Tunneling Protocol Parameters. This field is used to store parameters to configure the tunnel between the home IWF and the serving IWF. At a minimum, this field contains the type of tunneling protocol to use and any tunneling protocol specific parameters. For example, this field may contain the identifier for the tunneling protocol L2TP and any additional parameters required to configure the L2TP tunnel between the serving IWF and the home IWF.
4. Accounting Server Association. This field is used to store information needed by the home IWF to generate accounting data on behalf of the end system. It contains the name of the accounting protocol (e.g. RADIUS), the DNS name of the accounting server and additional parameters specific to the accounting protocol like the UDP port number, the shared secret that the IWF must use in the Radius Accounting protocol, the frequency of check-

US 6,421,714 B1

35

pointing, the seed for creating the session/multi-session id, etc. The accounting server's DNS name is translated to the accounting server's IP address, which is sent to the foreign agent.

For wireless service providers that have roaming agreements with each other, the FDD is used to do authentication and complete the registration process. If an end system roams from its home network to a foreign network, the registration server in the home network consults the FDD in its MSC to get information and authenticate the foreign network providing service to the end system.

The foreign domain directory management software provides a graphical user interface (GUI) based FDD management interface for system administrators. Using this GUI, system administrators are able to view and update entries in the FDD. This GUI is not intended for use by foreign wireless network service providers to perform remote updates based on roaming agreements. It is only intended for use by trusted personnel of the home wireless service provider operating behind firewalls.

The internet service provider directory (ISPD) is used by the home IWF to manage connectivity with ISPs that have service agreements with the wireless service provider so that subscribers may access their ISPs using the network. For each subscriber, the subscriber directory has an entry for the subscriber's ISP. This field points to an entry in the ISPD. The home IWF uses this information to set up the connection to the ISP on behalf of the subscriber.

The network architecture supports roaming. In order for roaming to work between wireless service providers, the architecture must support the setting up of roaming agreements between wireless service providers. This implies two things: (1) updating system directories across wireless service providers and (2) settlement of bills between service providers.

In order to allow subscribers access to internet service providers, the architecture supports roaming agreements with internet service providers. This implies that the architecture must be able to send data to and receive data from ISP PPP servers (i.e., that support industry standard protocols like PPP, L2TP and Radius). It also implies that the architecture handles directory updates for ISP access and settlement of bills with ISPs.

When roaming agreements are established between two wireless service providers, both providers have to update their home and foreign domain directories in order to support authentication and registration functions for end systems visiting their networks from the other network. At a minimum, the architecture of the present embodiment supports manual directory updates. When a roaming agreement is established between two wireless service providers, then the two parties to the agreement exchange information for populating their home and foreign domain directories. The actual updates of the directories is done manually by the personnel of the respective service providers. If later, the information in the home and foreign domain directories needs to be updated, the two parties to the agreement exchange the updated information and then manually apply their updates to the directories.

In an alternative embodiment, the directory management software incorporates developing standards in the IETF to enable roaming between internet service providers and to enable ISPs to automatically manage and discover roaming relationships. This makes manual directory management no longer necessary. The network system automatically propagates roaming relationships, and discovers them, to authenticate and register visiting end systems.

36

At a minimum, the network architecture just processes and stores the accounting data and makes the data available to the wireless service provider's billing system. It is up to the billing system to handle settlement of bills for roaming.

In an alternative embodiment, developing standards in the IETF to handle distribution of accounting records between internet service providers are incorporated into the network architecture to enable ISPs to do billing settlement for roaming end systems.

The system software supports access to ISPs and private intranets by supporting L2TP between the home IWF and the ISPs or intranet PPP server. The internet service provider directory contains information useful to the IWF for creating these tunnels. As access agreements between the wireless service provider and internet service providers are put in place, this directory is updated manually by the wireless service provider's personnel. Automatic updates and discovery of access relationships between the wireless service provider and internet service providers are presently contemplated and implemented as the internet standards evolve. While accessing an internet service provider, the subscriber receives two bills—one from the wireless service provider for the use of the wireless network and the second from the internet service provider. Although common billing that combines both types of charges is not handled by the minimum embodiment software, it is contemplated that the software will take advantage of internet standards for billing settlement as they evolve so that subscribers may receive a common bill based on roaming agreements between the ISP and wireless service providers.

The system includes a element management system for managing the network elements. From the element manager, system administrators perform configuration, performance and fault/alarm management functions. The element management applications run on top of a web browser. Using a web browser, system administrators manage the network from anywhere that they have TCP/IP access. The element manager also performs an agent role for a higher level manager. In this role it exports an SNMP MIB for alarm and fault monitoring.

A higher level SNMP manager is notified of alarm conditions via SNMP traps. The higher level SNMP manager periodically polls the element manager's MIB for the health and status of the network. System management personnel at the higher level manager are able to view an icon representation of the network and its current alarm state. By pointing and clicking on the network element icon, systems management personnel execute element management applications using a web browser and perform more detailed management functions.

Inside the network, management of the physical and logical network elements is performed using a combination of the SNMP protocol and internal management application programming interfaces. Applications in the element manager use SNMP or other management APIs to perform network management functions.

Architecturally, the element management system includes two distinct sets of functional elements. The first set of functional elements, including the configuration data server, performance data monitor and health/status monitor and network element recovery software, executes on an HA server equipped with RAID disks. The second set of functional elements, including the management applications, executes on a dedicated, non-HA management system. Even if the element manager system becomes non-operational, the network elements continue to be able to run and report alarms and even be able to recover from fault conditions.

US 6,421,714 B1

37

However, since all the management applications execute in the non-HA element manager, if the element manager goes down, then recovery actions requiring human intervention are not possible until the element manager becomes operational.

The wireless hubs (WIs) in the base stations are typically owned by a wireless service provider (WSP), and they are connected to the WSP's registration server (RS) either via point-to-point links, intranets or the Internet. The WSP's registration server is typically a software module executing on a processor to perform certain registration functions. Inter-working function units (IWF units) are typically software modules executing on a processor to perform certain interfacing functions. IWF units are typically connected to the registration servers via intranets/WAN, and the IWF units are typically owned by the WSP. However, the IWF units need not be located within the same LAN as the registration servers. Typically, accounting and directory servers (also software modules executing on a processor) are connected to the registration servers via a LAN in the service provider's Data Center (e.g., a center including one or more processors that hosts various servers and other software modules). Traffic from the end system is then routed via a router (connected to the LAN) to the public Internet or to an ISP's intranet.

The registration server located in a foreign WSP's network is referred to as the foreign registration server (FRS), and the registration server located in the end system's home network (where the mobile purchases its service) is referred to as the home registration server (HRS). The inter-working function unit in the home network is referred to as the home IWF while the inter-working function unit in the foreign network (i.e., the network the end system is visiting) is referred to as the serving IWF.

For fixed wireless service (i.e., a non-moving end system), an end system may register for service on the home network from the home network (e.g., at home service) or from a foreign network (e.g., roaming service). The end system receives an advertisement sent by an agent (e.g., an agent function implemented in software) in the wireless hub via the access point. There are both MAC-layer registration as well as network-layer registration to be accomplished. These may be combined together for efficiency.

For end systems at home (FIG. 23), the network layer registration (like a local registration) make's known to the home registration server the wireless hub to which the end system is currently attached. An IWF in the end system's home network will become the anchor or home IWF. Thus, PPP frames to and from the end system travel via the wireless hub to the home IWF in the home network. If the end system is at home, the home IWF is connected to the wireless hub via an XTunnel protocol.

For roaming wireless service (FIG. 24), the foreign registration server determines the identity of the home network of the roaming end system during the registration phase. Using this information, the foreign registration server communicates with the home registration server to authenticate and register the end system. The foreign registration server then assigns a serving IWF, and an I-XTunnel protocol connection is established between the home IWF and the serving IWF for the roaming end system. The serving IWF relays frames between the wireless hub and the home IWF. From the home IWF, data is sent to a PPP server (i.e., point-to-point protocol server) which may reside in the same IWF. However, if the data is to go to a corporate intranet or an ISP's intranet that has its own PPP server, the data is sent to the separate PPP server via the L2TP protocol. The

38

separate server is typically owned and operated by an Internet service provider who is different from the wireless service provider. For the duration of the session, the locations of the home IWF and PPP server remain fixed. The MAC layer registration can be combined with the network registration to economize on the overhead of separate communications for MAC layer and network layer registration; however, it may be advantageous to not combine these registration processes so that the WSP's equipment will be interoperable with other wireless networks that supports pure IETF Mobile-IP.

Registration sets up three tables. Table 1 is associated with each access point, and Table 1 identifies each connection (e.g., each end system) by a connection id (CID) and associates the connection id with a particular wireless (WM) modem address (i.e., the address of the end system or end system). Table 2 is associated with each wireless hub (WH), and Table 2 associates each connection id with a corresponding wireless modem address, access point and XTunnel id (XID). Table 3 is associated with each inter-working function (IWF), and Table 3 associates each connection id with a corresponding wireless modem address, wireless hub address, XTunnel id and IP port (IP/port). The entries described for these tables are described to include only relevant entries that support the discussion of mobility management. In reality, there are other important fields that need to be included as well.

TABLE 1

Connection Table at AP

CID	WM
C1	WM1
C2	WM1
C1	WM2

TABLE 2

Connection Table at WH

CID	WM	AP	XID
C1	WM1	AP1	5
C2	WM1	AP1	5
C1	WM2	AP1	6
C1	WM3	AP2	7

TABLE 3

Connection Table at IWF

CID	WM	WH	XID	IP/Port
C1	WM1	WH1	5	IP1/P1
C2	WM1	WH1	5	IP1/P2
C1	WM2	WH1	6	IP2/P3
C1	WM3	WH1	7	IP3/P1
C5	WM5	WH2	8	IP4/P1

The protocol stacks for dial-up users at home in a network as well as roaming users are illustrated in FIGS. 25-28. FIG. 25 depicts protocol stacks used for direct internet access by a fixed (i.e., non-moving) end system at home where a PPP protocol message terminates in the home IWF (typically collocated with the wireless hub) which relays message to and from an IP router and from there to the public internet. FIG. 26 depicts protocol stacks used for remote intranet access (i.e., either private corporate nets or an ISP) by a fixed

US 6,421,714 B1

39

(i.e., non-moving) end system at home where a PPP protocol message is relayed through the home IWF (typically collocated with the wireless hub) to a PPP server of the private corporate intranet or ISP. FIG. 27 depicts protocol stacks used for direct internet access by a roaming but fixed (i.e., non-moving) or a moving end system where the PPP protocol terminates in the home IWF (typically located in a mobile switching center of the home network) which relays message to and from an IP router. In FIG. 27, note how message traffic passes through a serving IWF (typically collocated with the wireless hub) in addition to the home IWF. FIG. 28 depicts protocol stacks used for remote intranet access (i.e., either private corporate nets or an ISP) by a roaming but fixed (i.e., non-moving) or a moving end system where a PPP protocol message is relayed through the home IWF (typically located in a mobile switching center of the home network) to a PPP server of the private corporate intranet or ISP. In FIG. 28, note how message traffic passes through a serving IWF (typically collocated with the wireless hub) in addition to the home IWF. When the serving IWF and the wireless hub are co-located in the same nest of computers or are even programmed into the same computer, it is not necessary to establish a tunnel using the X Tunnel protocol between the serving IWF and the wireless hub.

Equivalent variations to these protocol stacks (e.g. the RLP can be terminated at the wireless hub rather than at the

40

DIX ethernet can be used for encapsulating the XWD MAC primitives but the invention is not limited thereto. The ethernet frame format for XWD control frames is illustrated in FIG. 30. The ethernet header contains a destination address, a source address and an ethernet type field. The destination address field contains the ethernet address of the MAC entity to which the primitive is being sent. For MAC primitives invoked by the MAC user, this field will contain the ethernet address of the MAC user. For broadcast primitives, this address will be the ethernet broadcast address. The source address field contains the ethernet address of the MAC entity invoking the primitive. The ethernet type field contains the ethernet type for XWD. Possible values are XWD_Control for control frames and XWD_Data for data frames. These values must be different from all the ethernet type values that have been standardized and must be registered with the controlling authority.

The ethernet frame then has an XWD header field. The XWD header can be 16 bits long and will only be present for XWD control frames. The fields are illustrated in FIG. 31. The ethernet frame also contains a protocol header, a PPP payload field and a XWD MAC field. The header values for primitives using ethernet encapsulation are illustrated in Table 4 below.

Primitive Name	Destination Address	Source Address	Ethernet Type	XWD MAC Primitive
M_Discover.Req	Broadcast or unicast MAC Provider	MAC User	XWD_Control	0
M_Discover.Cnf	MAC User	MAC Provider	XWD_Control	1
M_OpenSap.Req	MAC Provider	MAC User	XWD_Control	2
M_OpenSap.Cnf	MAC User	MAC Provider	XWD_Control	3
M_CloseSap.Req	MAC Provider	MAC User	XWD_Control	4
M_CloseSap.Cnf	MAC User	MAC Provider	XWD_Control	5
M_EchoSap.Req	MAC User	MAC Provider	XWD_Control	6
M_EchoSap.Cnf	MAC Provider	MAC User	XWD_Control	7
M_Connect.Req	MAC Provider (modem only)	MAC User (end system only)	XWD_Control	8
M_Connect.Ind	MAC User (wireless hub only)	MAC Provider (AP only)	XWD_Control	9
M_Connect.Rsp	MAC Provider (AP only)	MAC User (wireless hub only)	XWD_Control	10
M_Connect.Cnf	MAC User (end system only)	MAC Provider (modem only)	XWD_Control	11
M_Disconnect.Req	MAC Provider	MAC User	XWD_Control	12

serving IWF or home IWF for mobiles at home) are also envisioned. If the IWF is located far from the wireless hub, and the packets can potentially be carried over a lossy IP network between the IWF and wireless hub, then it would be preferred to terminate the RLP protocol at the wireless hub. Another variation is the X tunnel between wireless hub and IWF need not be built on top of the UDP/IP. Xtunnels can be built using the Frame Relay/ATM link layer. However, the use of UDP/IP makes it easier to move the wireless hub and IWF software from one network to another.

Furthermore, the PPP protocol can be terminated in a wireless modem and sent to one or more endsystems via an ethernet connection. As illustrated in FIG. 29, the wireless modem 300 receives the PPP protocol information and encapsulates the PPP payload in an ethernet frame to be transported to at least one of the end systems 304 and 306 via the internet connection 302.

In another alternative, the PPP protocol can be terminated in a wireless router and then sent on to at least one end system connected to a local area network (LAN). As illustrated in FIG. 32, the wireless router 350 receives the PPP protocol information via the wireless modem 352. The router 354 receives the PPP information from the wireless modem 352 and sends the PPP information to at least one of the end systems 356, 358, 360 via a LAN link 362.

Four types of handoff scenarios may occur, and they are labeled: (i) local mobility, (ii) micro mobility, (iii) macro mobility, and (iv) global mobility. In all four scenarios (in one embodiment of the invention), a route optimization option is not considered so that the locations of the home registration server and the ISP's PPP server do not change. In another embodiment of the system with route optimization, the ISP's PPP server may change. However,

US 6,421,714 B1

41

this aspect is discussed below. In addition, the locations of the foreign registration server and IWF do not change in the first three scenarios.

The proposed IETF Mobile IP standard requires that whenever an end system changes the IP subnet to which it is attached, it sends a registration request message to a home agent in its home subnet. This message carries a care-of address where the end system can be reached in the new subnet. When traffic is sent, for example, from an ISP to an end system, the home agent intercepts the traffic that is bound for the end system as it arrives in the home subnet, and then forwards the traffic to the care-of address. The care-of address identifies a particular foreign agent in the foreign subnet. An end system's foreign agent can reside in the end system itself, or in a separate node that in turn forwards traffic to the end system (i.e., proxy registration agent). Mobile IP handoffs involve exchange of control messages between an end system's agent, the end system's home agent and potentially its corresponding hosts (CHs) (with route optimization option).

The proposed IETF Mobile IP standard would find it difficult to meet the latency and scalability goals for all movements in a large internetwork. However, the present hierarchical mobility management meets such goals. For small movements (e.g. a change of Access Points), only MAC-layer re-registrations are needed. For larger movements, network-layer re-registrations are performed. The present hierarchical mobility management is different from the flat-structure used in the IETF proposed Mobile-IP standard as well as the serving/anchor inter-working function model used in cellular systems like CDPD (based on a standard sponsored by the Cellular Digital Packet Data forum).

As depicted in FIG. 33, the local mobility handoff handles end system (designated MN for mobile node) movement between APs that belong to the same wireless hub. Thus, only MAC layer re-registration is required. The end system receives a wireless hub advertisement from a new AP and responds with a registration request addressed to the new AP.

The new AP (i.e., the one that receives the registration request from the end system) creates new entries in its connection table and relays the registration message to its wireless hub. In local mobility handoffs, the wireless hub does not change. The wireless hub recognizes the end system's registration request as a MAC level registration request, and it updates its connection table to reflect the connection to the new AP. Then, the old AP deletes the connection entry from its connection table. There are at least three ways whereby the old AP can delete the old entries, namely (i) upon time out, (ii) upon receiving a copy of the relayed MAC layer association message from the new AP to the wireless hub (if this relay message is a broadcast message), and (iii) upon being informed by the wireless hub of the need to delete the entry.

As depicted in FIG. 34, the micro mobility handoff handles end system (designated MN for mobile node) movement between wireless hubs that belong to the same registration server and where the end system can still be served by the existing serving IWF. When an advertisement is received from a new wireless hub (through a new AP), the end system sends a message to request registration to the registration server. The registration request is relayed from the new AP to the new wireless hub to the registration server.

When the registration server determines that the existing IWF can still be used, the registration server sends a build XTunnel Request message to request the existing IWF to build an XTunnel to the new wireless hub. Later, the

42

registration server sends a tear down XTunnel request message to request the existing IWF to tear down the existing XTunnel with the old wireless hub. The build and tear XTunnel Request messages can be combined into one message. A foreign registration server does not forward the registration message to the home registration server since there is no change of IWF, either the serving IWF or home IWF.

Upon receiving a positive build XTunnel reply and a positive tear XTunnel reply from IWF, the registration server sends a registration reply to end system. As the registration reply reaches the new wireless hub, the connection table at the new wireless hub is updated to reflect the connection to the new AP. The new AP updates its MAC filter address table and connection table after receiving a message from the new wireless hub, and the registration reply is forwarded to the end system.

The registration server sends a release message to the old wireless hub. When the old wireless hub receives the release message, it updates its connection table and the MAC filter address table and connection table of the old AP.

As depicted in FIG. 35, the macro mobility handoff case handles movement between wireless hubs that involves a change of the serving IWF in the foreign network, but it does not involve a change in the registration server. When an advertisement is received from a new wireless hub (through a new AP), the end system sends a message to request a network layer registration to the registration server. The registration request is relayed from the new AP to the new wireless hub to the registration server.

The registration server recognizes that it is a foreign registration server when the end system does not belong to the present registration server's network. This foreign registration server determines the identity of the home registration server by using a request, preferably a Radius Access request (RA request), to the foreign directory server (like a big yellow pages) and then assigns an appropriate IWF to be the serving IWF, and it forwards a registration request to the home registration server, preferably through a Radius Access request (RA request), informing the home registration server of the newly selected IWF.

The home registration server authenticates the registration request by using a request, preferably a Radius Access request (RA request), to the home directory server. Upon authenticating the request and determining that the existing home IWF can still be used, the home registration server instructs the home IWF to build a new I-XTunnel to the newly assigned serving IWF and to tear down the existing I-XTunnel to the old serving IWF. Upon receiving a positive build I-XTunnel reply and a positive tear I-XTunnel reply from the home IWF, the home registration server sends a registration reply to the foreign registration server.

The foreign registration server then instructs the newly assigned IWF to build an XTunnel to the new wireless hub. Upon receiving a positive build XTunnel reply, the foreign registration server instructs the old IWF to tear down the XTunnel to the old wireless hub. Upon receiving a positive build XTunnel reply and a positive tear XTunnel reply, the foreign registration server sends a registration reply to end system.

As the registration reply reaches the new wireless hub, the connection table at the new wireless hub is updated to reflect the connection to the new AP. The new AP updates its MAC filter address table and connection table after receiving a message from the new wireless hub, and the registration reply is forwarded to the end system.

The registration server sends a release message to the old wireless hub. When the old wireless hub receives the release

US 6,421,714 B1

43

message, it updates its connection table and the MAC filter address table, and the old AP updates its MAC filter address table and connection table after receiving a message from the old wireless hub.

The global mobility handoff case handles movement between wireless hubs that involves a change of registration servers. FIG. 36 depicts a global mobility handoff where the home IWF does not change, and FIG. 37 depicts a global mobility handoff where the home IWF changes. When an advertisement is received from a new wireless hub (through a new AP) in a new foreign network, the end system sends a message to request a network layer registration to the new foreign registration server. The registration request is relayed from the new AP to the new wireless hub to the new foreign registration server.

The registration server recognizes that it is a new foreign registration server when the end system does not belong to the present registration server's network. This foreign registration server determines the identity of the home registration server by using a request, preferably a Radius Access request (RA request), to the foreign directory server (like a big yellow pages) and then assigns an appropriate IWF to be the serving IWF, and it forwards the registration request to the home registration server, preferably through a Radius Access request (RA request), informing the home registration server of the newly selected IWF.

The home registration server authenticates the registration request by using a request, preferably a Radius Access request (RA request), to the home directory server. Upon authenticating the request and determining that the existing home IWF can still be used (FIG. 36), the home registration server instructs the home IWF to build a new I-XTunnel to the serving IWF newly assigned by the new foreign registration server. The home registration server also sends a de-registration message to the old foreign registration server and instructs the home IWF to tear down the existing I-XTunnel to the existing serving IWF of the old foreign network. Upon receiving a positive build I-XTunnel reply and a positive tear I-XTunnel reply from the home IWF, the home registration server sends a registration reply to the new foreign registration server.

The new foreign registration server then instructs the newly assigned IWF to build an XTunnel to the new wireless hub. Upon receiving a positive build XTunnel reply, the foreign registration server sends a registration reply to end system. As the registration reply reaches the new wireless hub, the connection table at the new wireless hub is updated to reflect the connection to the new AP. The new AP updates its MAC filter address table and connection table after receiving a message from the new wireless hub, and the registration reply is forwarded to the end system.

The old foreign registration server instructs the old IWF to tear down the XTunnel to the old wireless hub. Upon receiving a positive tear XTunnel reply or contemporaneously with the tear down XTunnel request, the old foreign registration server sends a release message to the old wireless hub. When the old wireless hub receives the release message, it updates its connection table and the MAC filter address table, and the old AP updates its MAC filter address table and connection table after receiving a message from the old wireless hub.

Alternatively, after the home registration server authenticates the registration request from the new foreign registration server and determines that the existing home IWF cannot be used (FIG. 37), the home registration server chooses a new home IWF and instructs the new home IWF to build a new level 2 tunnel protocol tunnel (L2TP tunnel)

44

to the present PPP server (e.g., the PPP server in a connected ISP intranet). Then, the home registration server instructs the old home IWF to transfer its L2TP tunnel traffic to the new home IWF.

Then the home registration server instructs the new home IWF to build a new I-XTunnel to the serving IWF newly assigned by the new foreign registration server. The home registration server also sends a de-registration message to the old foreign registration server and instructs the home IWF to tear down the existing I-XTunnel to the existing serving IWF of the old foreign network. Upon receiving a positive build I-XTunnel reply and a positive tear I-XTunnel reply from the home IWF, the home registration server sends a registration reply to the new foreign registration server.

The new foreign registration server then instructs the newly assigned IWF to build an XTunnel to the new wireless hub. Upon receiving a positive build XTunnel reply, the foreign registration server sends a registration reply to end system. As the registration reply reaches the new wireless hub, the connection table at the new wireless hub is updated to reflect the connection to the new AP. The new AP updates its MAC filter address table and connection table after receiving a message from the new wireless hub, and the registration reply is forwarded to the end system.

The old foreign registration server instructs the old IWF to tear down the XTunnel to the old wireless hub. Upon receiving a positive tear XTunnel reply or contemporaneously with the tear down XTunnel request, the old foreign registration server sends a release message to the old wireless hub. When the old wireless hub receives the release message, it updates its connection table and the MAC filter address table, and the old AP updates its MAC filter address table and connection table after receiving a message from the old wireless hub.

End systems constructed according to the present system interoperate with networks constructed according to the proposed IETF Mobile-IP standards, and end systems constructed according to the proposed IETF Mobile-IP standards interoperate with networks constructed according to the present invention.

Differences between the present system and the IETF Mobile-IP (RFC2002, a standards document) include:

- (i) The present systemists a hierarchical concept for mobility management rather than a flat structure as in the proposed IETF Mobile-IP standard. Small mobility within a small area does not result in a network level registration. Micro mobility involves setting up of a new Xtunnel and tearing down of an existing Xtunnel. Global mobility, at the minimum, involves setting up of a new I-XTunnel and tearing down of an existing I-XTunnel apart from the setting up/tearing down of XTunnel. Global mobility sometimes also involves setting up a new L2TP Tunnel and transferring of L2TP state from the existing L2TP Tunnel to the new L2TP Tunnel.
- (ii) In the present invention, a user name plus a realm is used to identify a remote dial-up user rather than a fixed home address as in the case of the proposed IETF Mobile-IP standard.
- (iii) In the present invention, registration and routing functions are carried out by separate entities. The two functions are carried out by the home agent in the proposed IETF Mobile IP standard, and both functions are carried out by the foreign agent in the proposed IETF Mobile IP standard. In contrast, in an embodiment of the present invention, registration is carried out in the registration server and routing functions are carried out by both the home and foreign IWF and the wireless hub (also referred to as the access hub).

US 6,421,714 B1

45

(iv) The present system utilizes three tunnels per PPP session. The XTunnel is more of a link-layer tunnel between the wireless hub and the serving IWF. The I-XTunnel between the serving IWF and the home IWF is more like the tunnel between home and foreign agents in the proposed IETF Mobile-IP standard. But it also has additional capabilities beyond the tunnels proposed by the Mobile-IP standard. The L2TP tunnel is used only when home IWF is not a PPP server. The number of these tunnels may be reduced by combining some functions in the same nodes as described earlier.

(v) In the present invention, wireless registration occurs before PPP session starts while in the proposed IETF Mobile-IP standard, Mobile-IP registration occurs after PPP session enters into the open state.

(vi) In the present invention, the network entity that advertises the agent advertisement (i.e., the wireless hub) is not on a direct link to the end systems whereas for the proposed IETF Mobile-IP standard, the agent advertisement must have a TTL of 1 which means that the end systems have a direct link with the foreign agent. In addition, the agent advertisement in the present systems not an extension to the ICMP router advertisements as in the proposed IETF Mobile-IP standard.

End systems in the present invention, should support agent solicitation. When an end system in the present system visits a network which is supporting the proposed IETF Mobile-IP standard, it waits until it hears an agent advertisement. If it does not receive an agent advertisement within a reasonable time frame, it broadcasts an agent solicitation.

In the present invention, network operators may negotiate with other networks that support the proposed IETF Mobile-IP standard such that home addresses can be assigned to the end systems of the present system that wish to use other networks. When the end system of the present system receives the agent advertisement, it can determine that the network it is visiting is not an a network according to the present system and hence uses the assigned home address to register.

For networks supporting the proposed IETF Mobile-IP standard, the PPP session starts before Mobile-IP registration, and the PPP server is assumed to be collocated with the foreign agent in such networks. In one embodiment, an SNAP header is used to encapsulate PPP frames in the MAC frames of the present system (in a manner similar to Ethernet format), and the foreign agent interprets this format as a proprietary PPP format over Ethernet encapsulation. Thus, the end system of the present system and its PPP peer can enter into an open state before the foreign agent starts transmitting an agent advertisement, and the end system of the present system can register.

To allow end systems supporting the proposed IETF Mobile-IP standard to work in networks of the type of the present invention, such mobiles are at least capable of performing similar MAC layer registrations. By making the agent advertisement message format similar to the proposed Mobile-IP standard agent advertisement message format, a visiting end system can interpret the agent advertisement and register with a wireless hub. In the present invention, registration request and reply messages are similar to the proposed IETF Mobile-IP standard registration request and reply messages (without any unnecessary extensions) so that the rest of the mobility management features of the present system are transparent to the visiting end systems.

Since end systems supporting the proposed IETF Mobile-IP standard expect a PPP session to start before Mobile-IP registration, an optional feature in wireless hubs of the

46

present system starts to interpret PPP LCP, NCP packets after MAC-layer registrations.

To avoid losing traffic during handoffs, the mobility management of the present systemists the make before break concept. For local mobility, a make before break connection is achieved by turning the MAC-layer registration message relayed by the new AP to the wireless hub into a broadcast message. That way, the old AP can hear about the new registration and forward packets destined for the end system that have not been transmitted to the new AP.

For micro mobility, information about the new wireless hub is included in the Tear XTunnel message exchanged between the serving IWF and the old WH. That way, the old wireless hub can forward buffered packets to the new wireless hub upon hearing a TearXTunnel message from the serving IWF. Alternatively, the RLP layer at the IWF knows the sequence number that has been acknowledged by the old wireless hub so far.

At the same time, the IWF knows the current send sequence number of the latest packet sent to the old wireless hub. Therefore, the IWF can forward those packets that are ordered in between these two numbers to the new wireless hub before sending newer packets to the new wireless hub. The RLP layer is assumed to be able to filter duplicate packet. The second approach is probably preferable to the first approach for the old wireless hub may not be able to communicate with one another directly.

For macro mobility, the old serving IWF can forward packets to the new serving IWF, in addition to the packet forwarding done from the old wireless hub to the new wireless. All we need to do is to forward the new serving IWF identity to the new serving IWF in the tear down I-XTunnel message. Another way to achieve the same result is to let the home IWF forward the missing packets to the new serving IWF rather than asking the old serving IWF to do the job since the home IWF knows the I-XTunnel sequence number last acknowledged by the old serving IWF and the current I-XTunnel sequence number sent by the home IWF.

The method of estimating how much buffer should be allocated per mobile per AP per wireless hub per IWF such that the traffic loss between handoffs can be minimized is to let the end system for the AP for the wireless hub for the IWF estimate the packet arrival rate and the handoff time. This information is passed to the old AP of the wireless hub of the IWF to determine how much traffic should be transferred to the new AP of the wireless hub of the IWF, respectively, upon handoffs.

To achieve route optimization in the present invention, the end system chooses the PPP server closest to the serving IWF. Without route optimization, excessive transport delays and physical line usage may be experienced.

For example, an end system subscribed to a home network in New York City may roam to Hong Kong. To establish a link to a Hong Kong ISP, the end system would have a serving IWF established in a wireless hub in Hong Kong and a home IWF established in the home network in New York City. A message would then be routed from the end system (roamed to Hong Kong) through the serving IWF (in Hong Kong) and through the home IWF (in New York City) and back to the Hong Kong ISP.

A preferred approach is to connect from the serving IWF (in Hong Kong) directly to the Hong Kong ISP. The serving IWF acts like the home IWF. In this embodiment, roaming agreements exist between the home and foreign wireless providers. In addition, the various accounting/billing systems communicate with one another automatically such that

US 6,421,714 B1

47

billing information is shared. Accounting and billing information exchange may be implemented using standards such as the standard proposed by the ROAMOPS working group of the IETF.

However, the serving IWF must still discover the closest PPP server (e.g., the Hong Kong ISP). In the present embodiment, the foreign registration server learns of the end system's desire to connect to a PPP server (e.g., a Hong Kong ISP) when it receives a registration request from the end system. When the foreign registration server determines that the serving IWF is closer to the desired PPP server (e.g., the Hong Kong ISP) than the home IWF is, the foreign registration server instructs the serving IWF to establish an L2TP tunnel to its nearest PPP server (in contrast to the PPP server closest to the home registration server and home IWF). Then, the foreign registration server informs the home registration server that the end system is being served by the serving IWF and the foreign PPP.

In an alternative embodiment, the foreign registration server determines that the serving IWF is closer to the desired PPP server (e.g., the Hong Kong ISP) than the home IWF is, when it receives a registration request from the end system. The foreign registration server relays the registration request message to the home registration server with an attached message indicating the serving IWF information and a notification that route optimization is preferred. At the same time, the foreign registration server instructs the serving IWF to establish an L2TP tunnel to the PPP server. Upon approving the registration request, the home registration server instructs the home IWF to transfer the L2TP state to the foreign IWF.

Having described preferred embodiments of a novel network architecture with wireless end users able to roam (which are intended to be illustrative and not limiting), it is noted that modifications and variations can be made by persons skilled in the art in light of the above teachings. For example, connection links described herein may make reference to known connection protocols (e.g., IP, TCP/IP, L2TP, IEEE 802.3, etc.); however, the system contemplates other connection protocols in the connections links that provide the same or similar data delivery capabilities. Acting agents in the above described embodiments may be in the form of software controlled processors or may be other form of controls (e.g., programmable logic arrays, etc.). Acting agents may be grouped as described above or grouped otherwise in keeping with the connection teachings described herein and subject to security and authentication teachings as described herein. Furthermore, a single access point, access hub (i.e., wireless hub) or inter-working function unit (IWF unit) may provide multi-channel capability. Thus, a single access point or access hub or IWF unit may act on traffic from multiple end systems, and what is described herein as separate access points, access hubs or IWF units contemplates equivalence with a single multi-channel access point, access hub or IWF unit. It is therefore to be understood that changes may be made in the particular embodiments of the system disclosed which are within the scope and spirit of the systems defined by the appended claims.

Having thus described the system with the details and particularity required by the patent laws, what is claimed and desired protected by Letters Patent is set forth in the appended claims.

What is claimed is:

1. A wireless data network comprising:

a home network that includes a home mobile switching center and a wireless end system, the home mobile

48

switching center including a home registration server and a home inter-working function, the wireless end system including an end registration agent, the end registration agent being coupled to the home registration server;

a PPP server, a message being coupleable from the end system through the home inter-working function to the PPP server;

a foreign network that includes a foreign base station with a foreign access hub, the foreign access hub including a first serving inter-working function; and

a second mobile end system subscribed to the home network and operating within the foreign network, a first message being transportable between the first mobile end system and a first communications server through the first home inter-working function and through the first serving inter-working function of the foreign access hub in the foreign base station.

2. The network of claim 1, wherein the first message is transportable from the second mobile end system through the first home inter-working function to the first communications server.

3. The network of claim 1, wherein the second mobile end system includes a wireless modem coupleable to the foreign access hub.

4. The network of claim 1, further comprising:

a third end system subscribed to the home network and operating as a fixed end system within the home network; and

a home base station that includes a home access hub with a second home inter-working function, a second message being transportable between the third end system and a second communications server through the second home inter-working function.

5. The network of claim 1, further comprising:

a third end system subscribed to the home network and operating as a mobile end system within the home network;

a home mobile switching center having a second home inter-working function; and

a home base station that includes a home access hub with a second serving inter-working function, a second message being transportable between the second end system and a second communications server through the second serving inter-working function and through the second home inter-working function.

6. The network of claim 1, wherein the first home inter-working function includes a home accounting collection module to collect accounting data on message traffic transported through the first home inter-working function.

7. The network of claim 6, wherein:

the home mobile switching center includes a home accounting server; and

the home accounting collection module includes a sub-module to periodically send accounting reports to a home accounting server.

8. The network of claim 7, wherein:

the home network further includes a home billing processor; and

the home accounting server includes a module to send accounting reports to the home billing processor, the home billing processor including a module to prepare customer bills based on the accounting reports from the home accounting server.

US 6,421,714 B1

49

9. The network of claim 8, wherein:

the home network further includes a home billing processor;

the foreign network further includes a foreign accounting server and a foreign billing processor;

the first serving inter-working function includes a foreign accounting collection module to collect accounting data on message traffic transported through the first serving inter-working function, the foreign accounting collection module including a sub-module to periodically send accounting reports to the foreign accounting server, the foreign accounting server including a module to send accounting reports to the foreign billing processor, the foreign billing processor including a module to send accounting reports to the home billing processor, the home billing processor including a module to prepare customer bills based on the accounting reports from the foreign billing processor.

10. A wireless data network comprising:

a home network that includes a home mobile switching center and a wireless end system, the home mobile switching center including a home registration server and a home inter-working function, the wireless end system including an end registration agent, the end registration agent being coupled to the home registration server;

a PPP server, a message being coupleable from the end system through the home inter-working function to the PPP server;

a foreign network that includes a foreign mobile switching center and a base station, the foreign mobile switching center including a serving registration server, the base station including an access hub, the access hub including a proxy registration agent; and

a second end system subscribed to the home network and operating within the foreign network, the second end system including an end registration agent, the end registration agent being coupled to the proxy registration agent, the proxy registration agent being coupled to the serving registration server, the serving registration server being coupled to the home registration server,

wherein the home registration server includes a module to authenticate that the foreign network is authorized to host the end system;

wherein the home registration server includes a module to authenticate that the end system is authorized to receive services of the home network;

wherein the serving registration server includes a module to authenticate that the end system is a subscriber of the home network;

wherein the home network further includes a home billing processor;

wherein the foreign network further includes a foreign accounting server and a foreign billing processor; and

wherein the first serving inter-working function includes a foreign accounting collection module to collect accounting data on message traffic transported through the first serving inter-working function, the foreign accounting collection module including a sub-module to periodically send accounting reports to the foreign accounting server, the foreign accounting server including a module to send accounting reports to the foreign billing processor, the foreign billing processor including a module to send accounting reports to the

50

home billing processor, the home billing processor including a module to prepare customer bills based on the accounting reports from the foreign billing processor.

11. A wireless data network comprising:

a home network that includes a home mobile switching center and a wireless end system, the home mobile switching center including a home registration server and a home inter-working function, the wireless end system including an end registration agent, the end registration agent being coupled to the home registration server;

a PPP server, a message being coupleable from the end system through the home inter-working function to the PPP server;

a foreign network that includes a base station with an access hub and a foreign mobile switching center with a serving registration server, the access hub including a serving inter-working function, the serving inter-working function including a foreign accounting collection module;

the home inter-working function including a home accounting collection module, and

a second end system subscribed to the wireless data network and coupled to the foreign access hub, the home and server accounting collection modules collecting accounting data on message traffic transported between the second end system and a communications server through the home inter-working function and through the serving inter-working function.

12. The network of claim 11, wherein the home accounting collection module includes a sub-module to collect accounting data on message traffic transported from the second end system through the home inter-working function to the communications server.

13. The network of claim 11, wherein:

the foreign mobile switching center includes a foreign accounting server; and

the foreign accounting collection module includes a sub-module to periodically send accounting reports to the foreign accounting server.

14. The network of claim 11, wherein:

the home mobile switching center includes a home accounting server;

the home accounting collection module includes a sub-module to periodically send accounting reports to the home accounting server.

15. The network of claim 14, wherein:

the home network further includes a home billing processor; and

the home accounting server includes a module to send accounting reports to the home billing processor, the home billing processor including a module to prepare customer bills based on the accounting reports from the home accounting server.

16. The network of claim 11, wherein:

the foreign network further includes a foreign accounting server and a foreign billing processor;

the foreign accounting collection module includes a sub-module to collect accounting data on message traffic transported through the first serving inter-working function, the foreign accounting collection module further including a sub-module to periodically send accounting reports to the foreign accounting server, the foreign accounting server including a module to send

US 6,421,714 B1

51

accounting reports to the foreign billing processor, the foreign billing processor including a module to send accounting reports to the home billing processor, the home billing processor including a module to prepare customer bills based on the accounting reports from the foreign billing processor.

17. The network of claim 11, wherein:

the foreign mobile switching center includes a foreign accounting server;

the home mobile switching center includes a home accounting server;

the foreign accounting collection module includes a sub-module to periodically send accounting reports to a foreign accounting server; and

the home accounting collection module includes a sub-module to periodically send accounting reports to a home accounting server.

18. A wireless data network comprising:

a home network that includes a home mobile switching center and a wireless end system, the home mobile switching center including a home registration server and a home inter-working function, the wireless end system including an end registration agent, the end registration agent being coupled to the home registration server;

a PPP server, a message being coupleable from the end system through the home inter-working function to the PPP server;

a foreign network that includes a foreign mobile switching center with a serving registration server;

the home mobile switching center includes a plurality of unassigned home inter-working functions; and

a second end system subscribed to the home network and operating within the foreign network, the second end system including an end registration agent to form a registration request, the end registration agent sending the registration request through the serving registration server to the home registration server, the home registration server including a module to select an active home inter-working function from the plurality of unassigned home inter-working functions based on the registration request.

19. The network of claim 18, wherein:

the serving inter-working function is regarded as an active serving inter-working function;

the foreign network further includes a plurality of serving inter-working functions; and

the serving registration server includes a module to select the active serving inter-working function from the plurality of serving inter-working functions based on the registration request.

20. The network of claim 18, wherein the home registration server includes a module to authenticate that the foreign network is authorized to host the second end system.

52

21. The network of claim 18, wherein the home registration server includes a module to authenticate that the second end system is authorized to receive services of the home network.

22. The network of claim 18, wherein the serving registration server includes a module to authenticate that the second end system is a subscriber of the home network.

23. The network of claim 18, wherein:

the registration request includes service type information; and

the home registration server includes a module to control the selection of the active home inter-working function based on the service type information.

24. The network of claim 23, wherein the service type information specifies a request for one of public internet service and private intranet service.

25. The network of claim 23, wherein the service type information specifies a request for one of mobile service and fixed service.

26. The network of claim 18, wherein:

the registration request includes quality of service information; and

the home registration server includes a module to control the selection of the active home inter-working function based on the quality of service information.

27. The network of claim 26, wherein the quality of service information specifies a request for one of constant bit rate service, real time variable bit rate service, non-real time variable bit rate service, unspecified bit rate service and available bit rate service.

28. A wireless data network comprising:

a home network that includes a home mobile switching center and a wireless end system, the home mobile switching center including a home registration server and a home inter-working function, the wireless end system including an end registration agent, the end registration agent being coupled to the home registration server;

a PPP server, a message being coupleable from the end system through the home inter-working function to the PPP server;

a foreign network that includes a foreign mobile switching center and a base station, the base station including an access hub with a serving inter-working function; and

a roaming end system subscribed to the home network and operating within the foreign network, a message being transportable between the roaming end system and the home inter-working function through the serving inter-working function using a protocol that ensures in sequence delivery of data packets.

* * * * *

EXHIBIT 14



US006493679B1

(12) **United States Patent**
Rappaport et al.

(10) **Patent No.: US 6,493,679 B1**
(45) **Date of Patent: Dec. 10, 2002**

(54) **METHOD AND SYSTEM FOR MANAGING A REAL TIME BILL OF MATERIALS**

(75) Inventors: **Theodore S. Rappaport**, Blacksburg, VA (US); **Roger R. Skidmore**, Blacksburg, VA (US)

(73) Assignee: **Wireless Valley Communications, Inc.**, Austin, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/318,842**

(22) Filed: **May 26, 1999**

(51) Int. Cl.⁷ **G06F 17/60**

(52) U.S. Cl. **705/29; 705/1; 705/28; 705/29; 705/10**

(58) Field of Search **705/28, 29, 1, 705/10; 702/184; 703/1, 6, 7, 13, 14, 17, 25, 22, 27, 28**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,111,392	A *	5/1992	Malin	705/29
5,119,307	A *	6/1992	Blaha et al.	705/29
5,239,487	A *	8/1993	Horejsi et al.	702/184
5,307,261	A *	4/1994	Maki et al.	705/29
5,515,269	A *	5/1996	Willis et al.	705/29
5,761,093	A *	6/1998	Mullen et al.	705/29
5,815,395	A *	9/1998	Hart et al.	705/29

FOREIGN PATENT DOCUMENTS

WO	WO 200013115	*	3/2000
WO	WO 200073953	*	12/2000
WO	WO 200074401	*	12/2000

OTHER PUBLICATIONS

"CoCreate Announces New Release of SolidDesigner; First System Enabling Users to Attach Parameters to Imported CAD Models", Business Wire, p0286, Apr. 7, 1999.*

Smith, Alan, "Speedy MicroStation release does 95", Machine Design, v68, n10, p98 (1), May 23, 1996.*

"Latest new products. (New Briefs)", Communications News, 37, 5, 98, May 2000.*

Tolputt, Elvadia, "ICL moves to consolidate its hold on the manufacturing market with OMAC 2000 under VME", Computergram International, n1371, CGI02230010, Feb. 23, 1990.*

Trappey et al, "Object-oriented bill of materials system for dynamic product management", Journal of Intelligent Manufacturing, v7, n5, Oct. 1996, p365-371.*

Gonzales, Daniel M., "ArchiCAD 4.02", Macworld, Jun. 1, 1992, v9, n6, p199, 201, 2 Pages.*

(List continued on next page.)

Primary Examiner—Vincent Millin

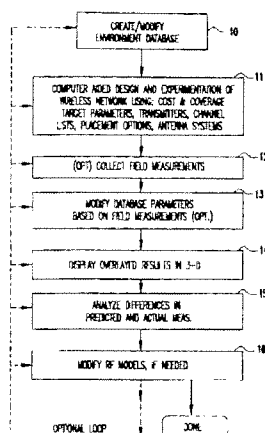
Assistant Examiner—Nga Nguyen

(74) *Attorney, Agent, or Firm*—Whitham, Curtis & Christofferson, PC

(57) **ABSTRACT**

An automated method for quickly generating a complete bill of materials and total cost information in real time. Components for a desired system are specified and/or replaced by substitute components, while continuously predicting the wireless system performance. A design engineer builds a model of the desired wireless communications system and specifies each component necessary to provide sufficient or optimal system performance. A parts list is maintained, in real time, that contains a definition of each system component and its associated performance and cost parameters. As the user changes wireless system designs through a series of "what-if" scenarios, components are replaced with substitute components, cable lengths are modified, antenna systems and base station parameters are re-designed and moved to alternate locations, etc. The bill of materials is automatically updated and component costs and total system costs are immediately available to the design engineer. The designer may choose to swap components for less expensive components or may investigate several alternate radio frequency distribution and antenna schemes, etc. The performance characteristics of the system are automatically updated as is the system cost as the designer to assesses the trade-offs between performance and cost at the same time.

19 Claims, 19 Drawing Sheets



US 6,493,679 B1

Page 2

OTHER PUBLICATIONS

- Bell Laboratories, "Wise—A Wireless System Engineering Tool", pub. date unknown, web pages from <http://www.bell-labs.com>.
- Mobile Systems International, "Planet", pub. date unknown, web pages from <http://www.msi-swe.se>.
- EDX Emgineerong. Inc., "EDX Engineering—Products: Version 2.0 of EDX SignalPro(tm)", pub. date unknown, web pages from <http://www.edx.com>.
- Ericsson Wireless Communications, "TEMS Total—TEMS Product Overview", pub. date unknown, web pages from <http://www.ericsson.com>.
- Ericsson Wireless Communications, "TEMS Total—TEMS Frequently asked questions—general", pub. date unknown, web pages from <http://www.ericsson.com>.
- Ericsson Wireless Communications, "TEMS Total—TEMS Light Frequently asked questions—general", pub. date unknown, web pages from <http://www.ericsson.com>.
- Ericsson Wireless Communications, "TEMS Total—TEMS Makes Your Net Work", pub. date unknown, web pages from <http://www.ericsson.com>.
- D. Ullmo et al., "Wireless Propagation in Buildings: A Statistical Scattering Approach", IEEE Transactions on Vehicular Technology, vol. 48, No. 3, May 1999.
- S. Seidel et al., "Site-Specific Propagation Prediction for Wireless In-Building Personal Communication System Design", IEEE Transactions on Vehicular Technology, vol. 43, No. 4, Nov. 1994.
- M. Panjwani et al., "Interactive Computation of Coverage Regions for Wireless Communication in Multifloored Indoor Environments", IEEE Journal on Selected Areas in Communications, vol. 14, No. 3, Apr. 1996.
- C. M. Peter Ho, et al., "Antenna Effects on Indoor Obstructed Wireless Channels and a Deterministic Image-Based Wide-Band Propagation Model for In-Building Personal Communication Systems", International Journal of Wireless Information Networks, vol. 1, No. 1, 1994.
- S. Kim et al., "Radio Propagation Measurements and Prediction Using Three-Dimensional Ray Tracing in Urban Environments at 908 MHz and 1.9 GHz", IEEE Transactions on Vehicular Technology, vol. 48, No. 3, May 1999.
- L. Piazzi et al., "Achievable Accuracy of Site-Specific Path-Loss Predictions in Residential Environments", IEEE Transactions on Vehicular Technology, vol. 48, No. 3, May 1999.
- S. Sandip et al., "Predicted Path Loss for Rosslyn, VA", The Bradley Department of Electrical Engineering, MPRG TR-95-03, Virginia Tech, Mar. 5, 1995.
- S. Sandip et al., "Predicted Path Loss for Rosslyn, VA", The Bradley Department of Electrical Engineering, First Set of Predictions, MPRG TR-94-20, Virginia Tech, Dec. 9, 1994.
- G. Durgin, et al., "Measurements and Models for Radio Path Loss and Penetration Loss in and Around Homes and Trees at 5.85 GHz", IEEE Transactions on Communications, vol. 46, No. 11, Nov. 1998.
- R. Skidmore et al., "A Comprehensive In-Building and Microcellular Wireless Communication System Design Tool", The Bradley Department of Electrical Engineering, MPRG-TR-97-13, Jun. 9, 1997.
- T. Rappaport, et al., "Radio Propagation Prediction Techniques and Computer-Aided Channel Modeling for Embedded Wireless Microsystems", The Bradley Department of Electrical Engineering, ARPA Annual Report, Jul. 1994.
- T. Rappaport, et al., "Radio Propagation Prediction Techniques and Computer-Aided Channel Modeling for Embedded Wireless Microsystems", The Bradley Department of Electrical Engineering, ARPA Semi-Annual Report, Jul. 1995.
- T. Rappaport et al., "Use of Topographic Maps With Building Information to Determine Antenna Placements and GPS Satellite Coverage for radio Detection and Tracking in Urban Environments", The Bradley Department of Electrical Engineering, Quarterly Report, Sep. 1995.
- T. Rappaport et al., "Use of Topographic Maps With Building Information to Determine Antenna Placements and GPS Satellite Coverage for radio Detection and Tracking in Urban Environments", The Bradley Department of Electrical Engineering, MPRG-TR-95-19, Nov. 1995.
- R. Skidmore et al., "Interactive Coverage Region and System Design Simulation for Wireless Communication Systems in Multi-floored Indoor Environments, SMT Plus" IEEE ICUPC '96 Proceedings.
- "Site Planner 3.16 for Windows 95/98NT User's Manual" (Wireless Valley Communications, Inc. 1999), pp. 5-148 to 5-156.

* cited by examiner

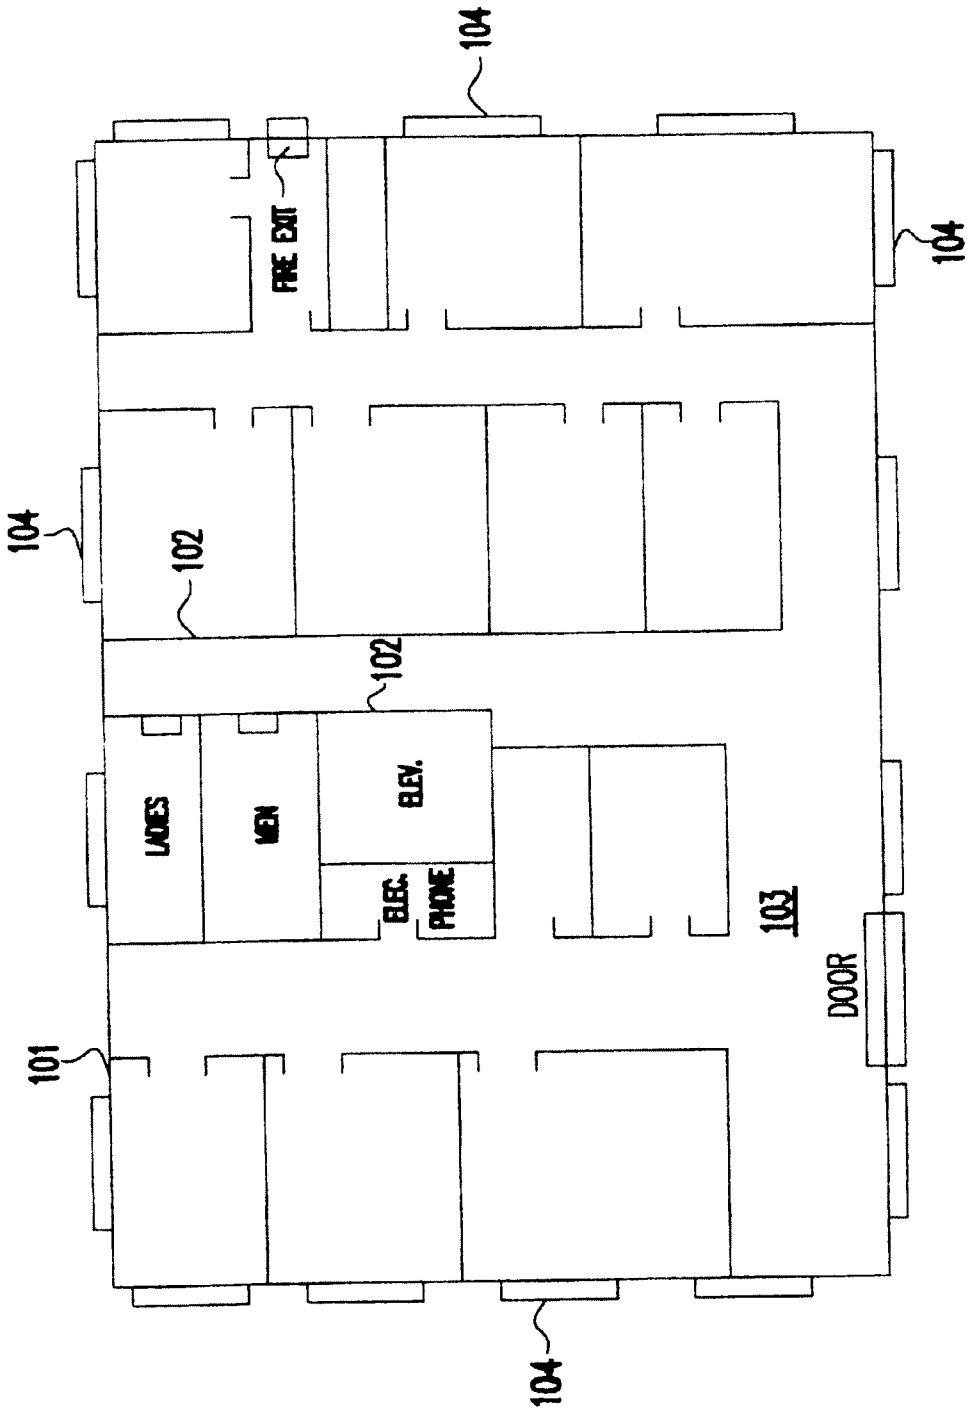
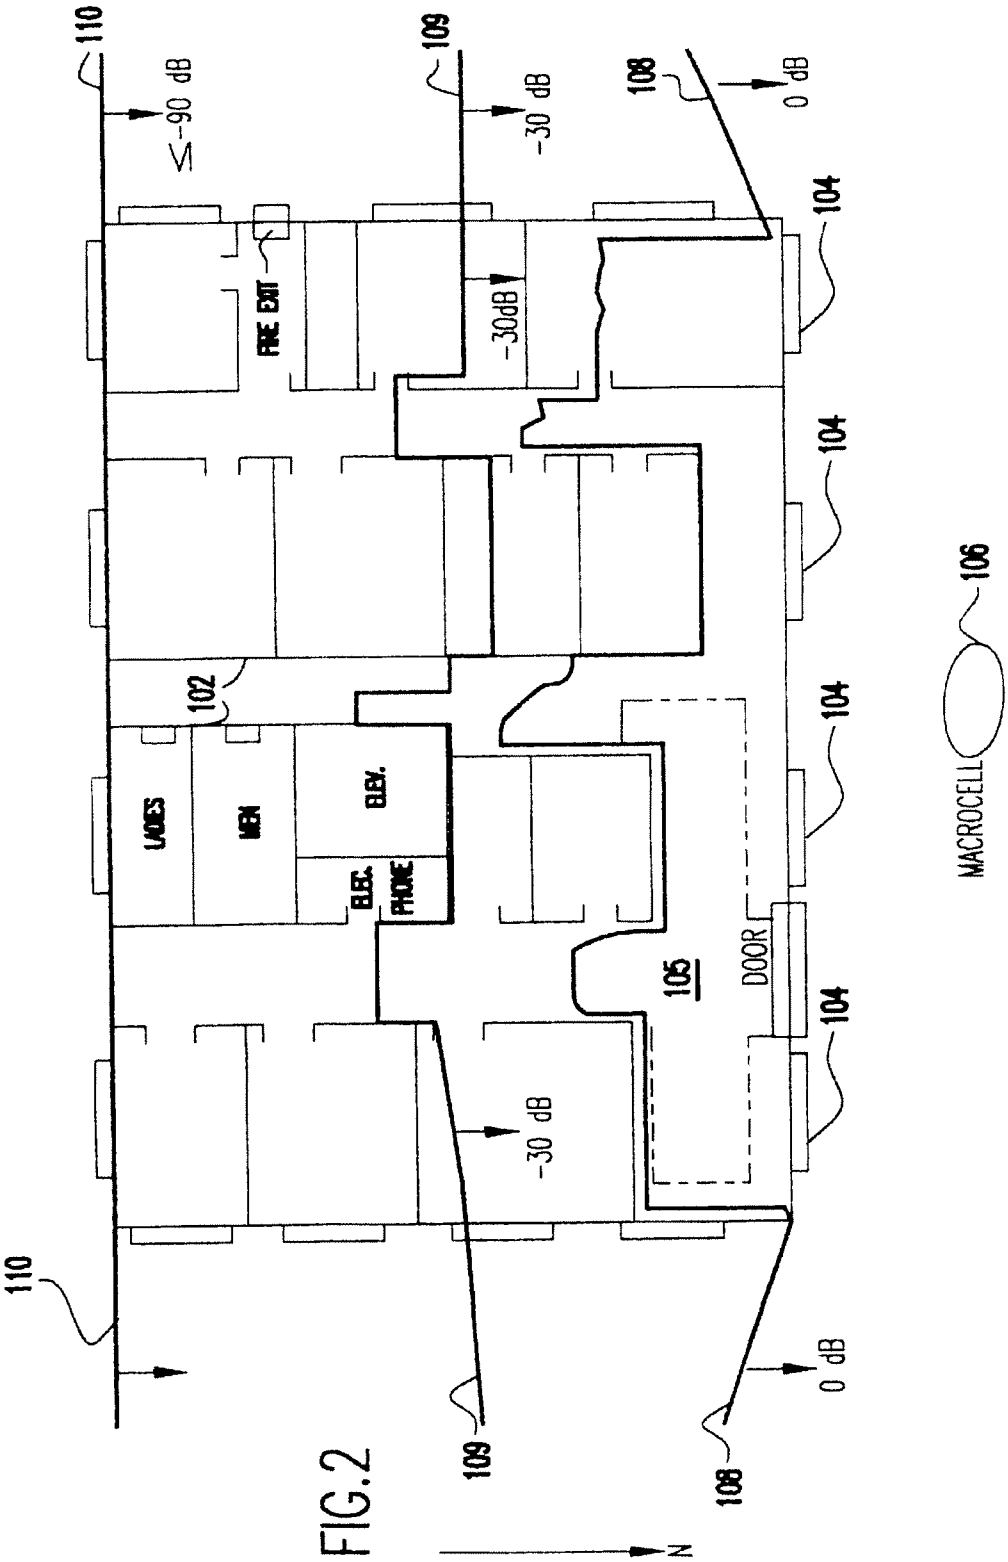


FIG. 1



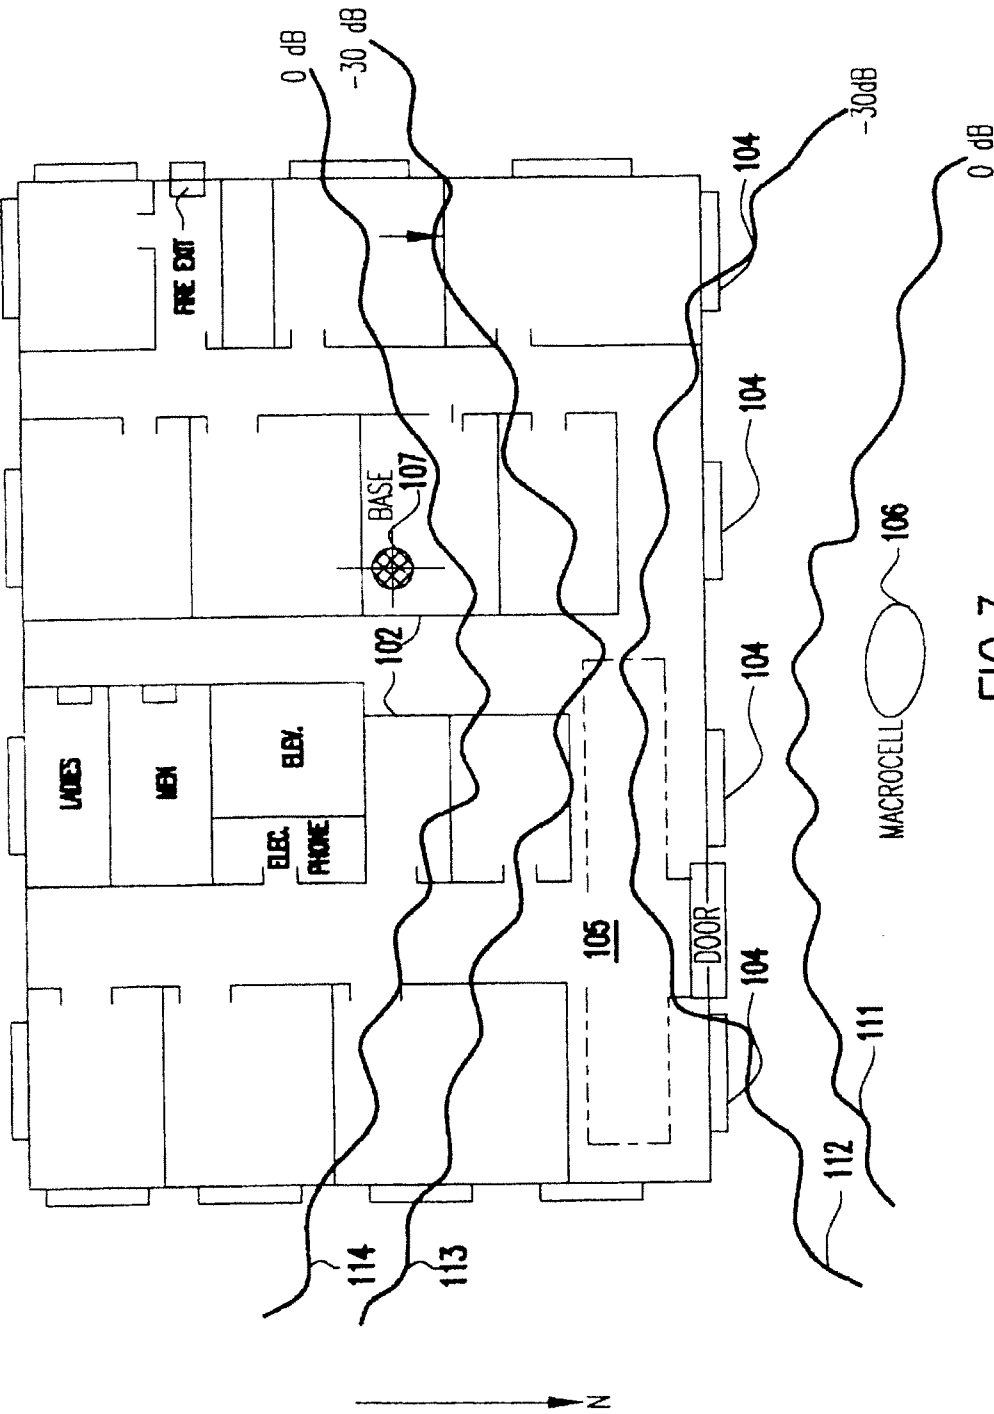


FIG. 3

U.S. Patent

Dec. 10, 2002

Sheet 4 of 19

US 6,493,679 B1

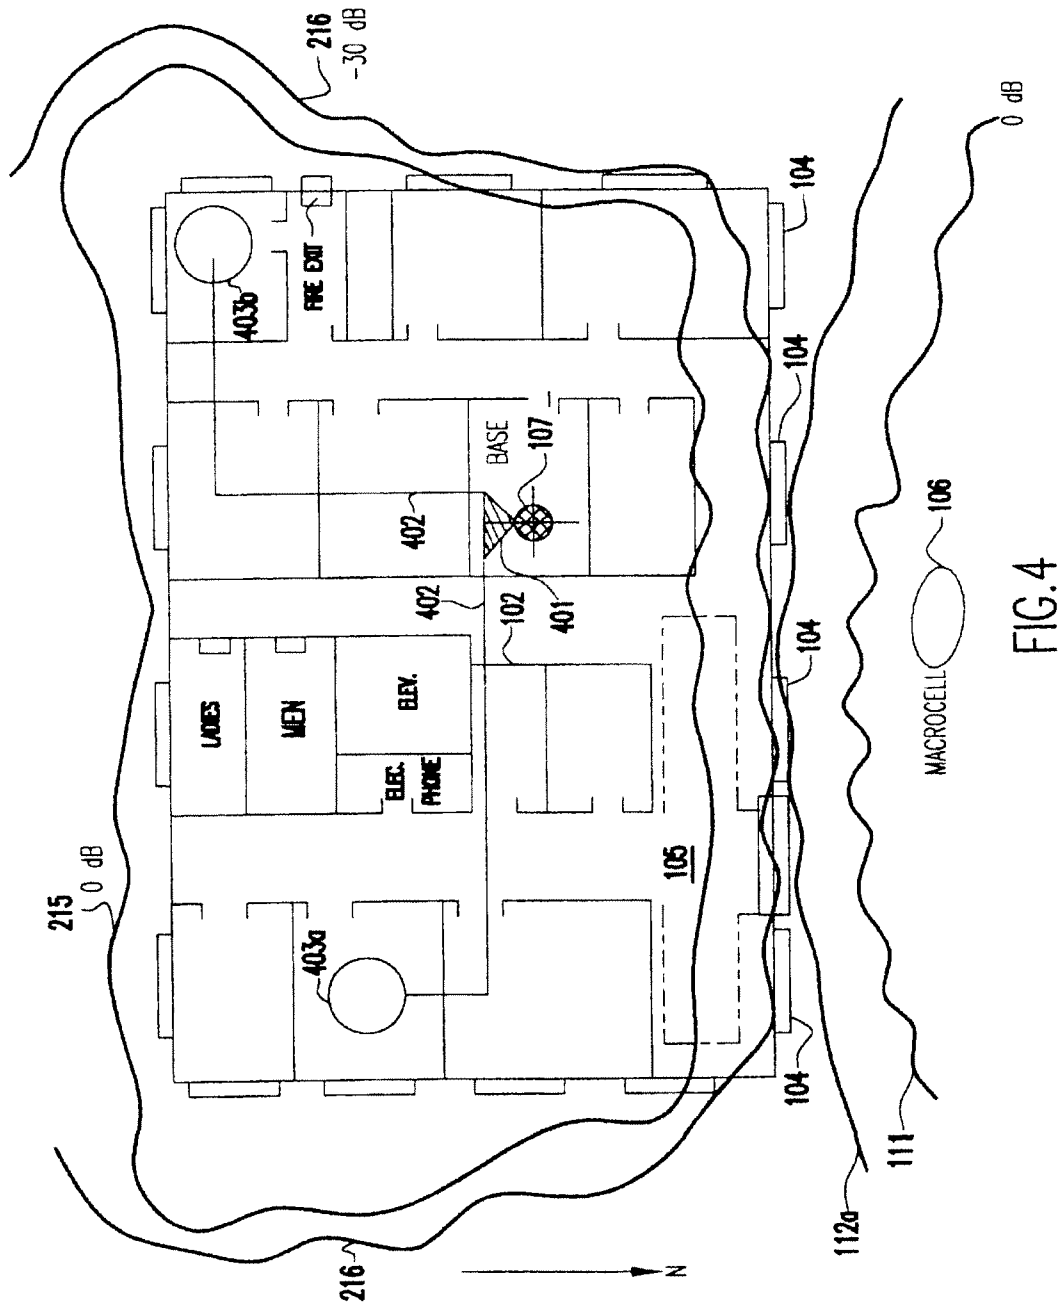


FIG. 4

U.S. Patent

Dec. 10, 2002

Sheet 5 of 19

US 6,493,679 B1

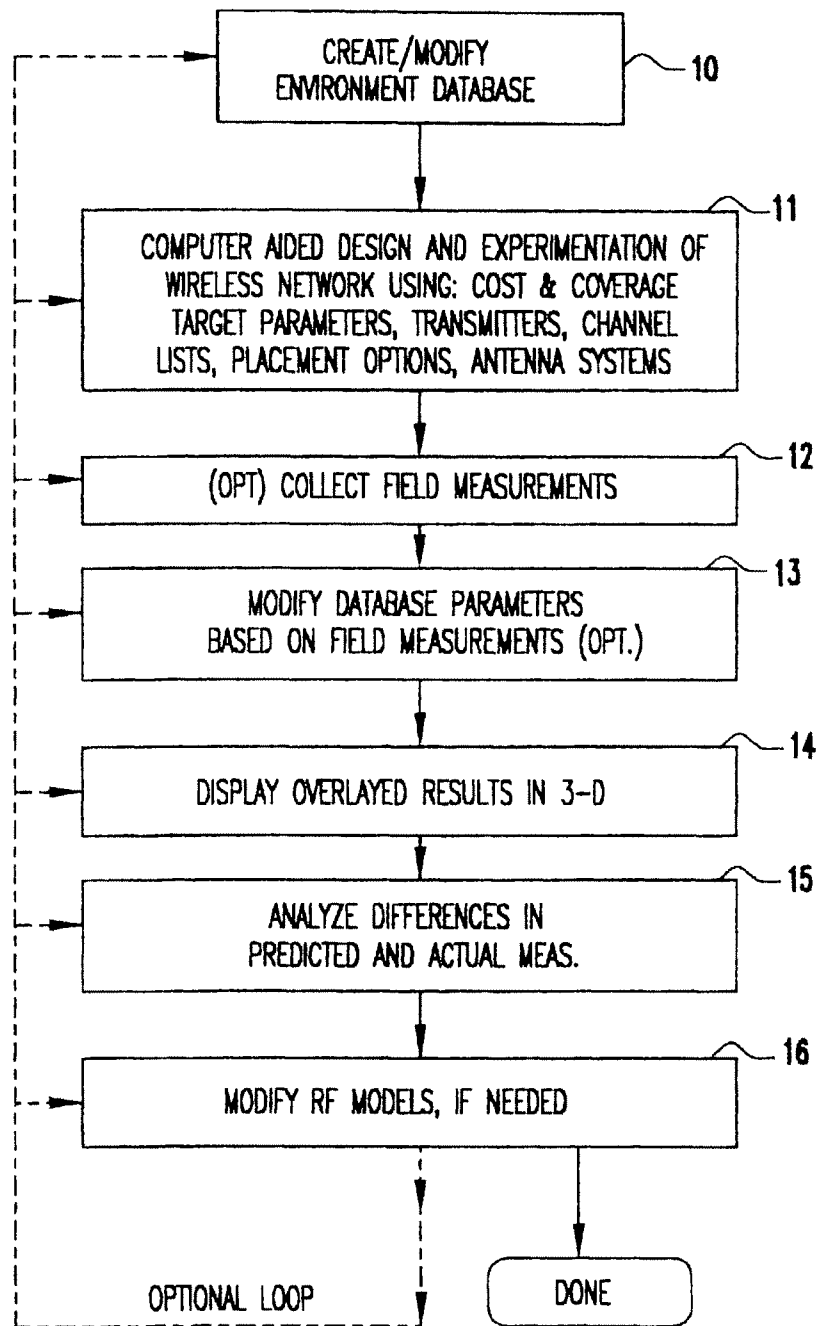


FIG.5

U.S. Patent

Dec. 10, 2002

Sheet 6 of 19

US 6,493,679 B1

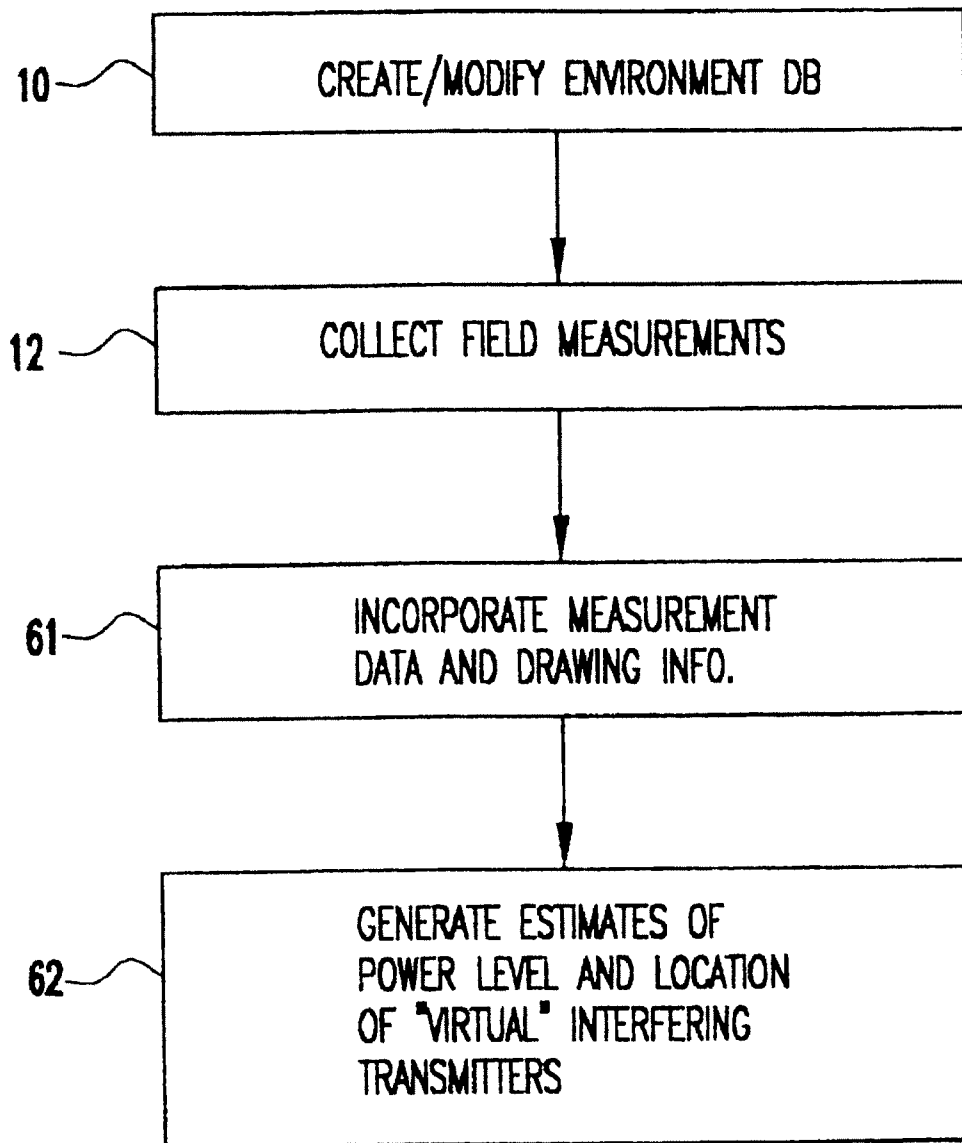


FIG.6

U.S. Patent

Dec. 10, 2002

Sheet 7 of 19

US 6,493,679 B1

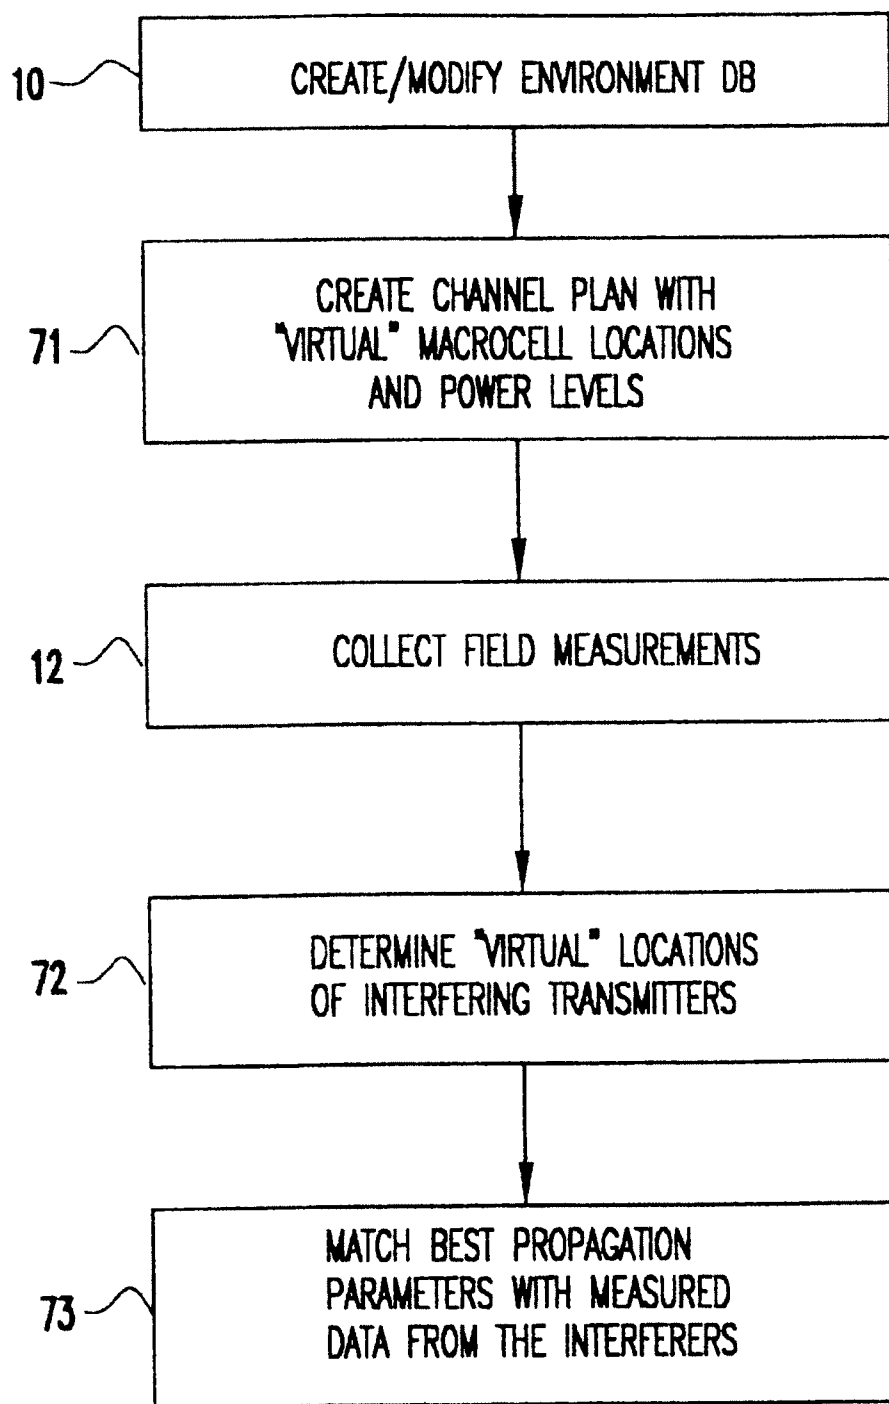


FIG.7

U.S. Patent

Dec. 10, 2002

Sheet 8 of 19

US 6,493,679 B1

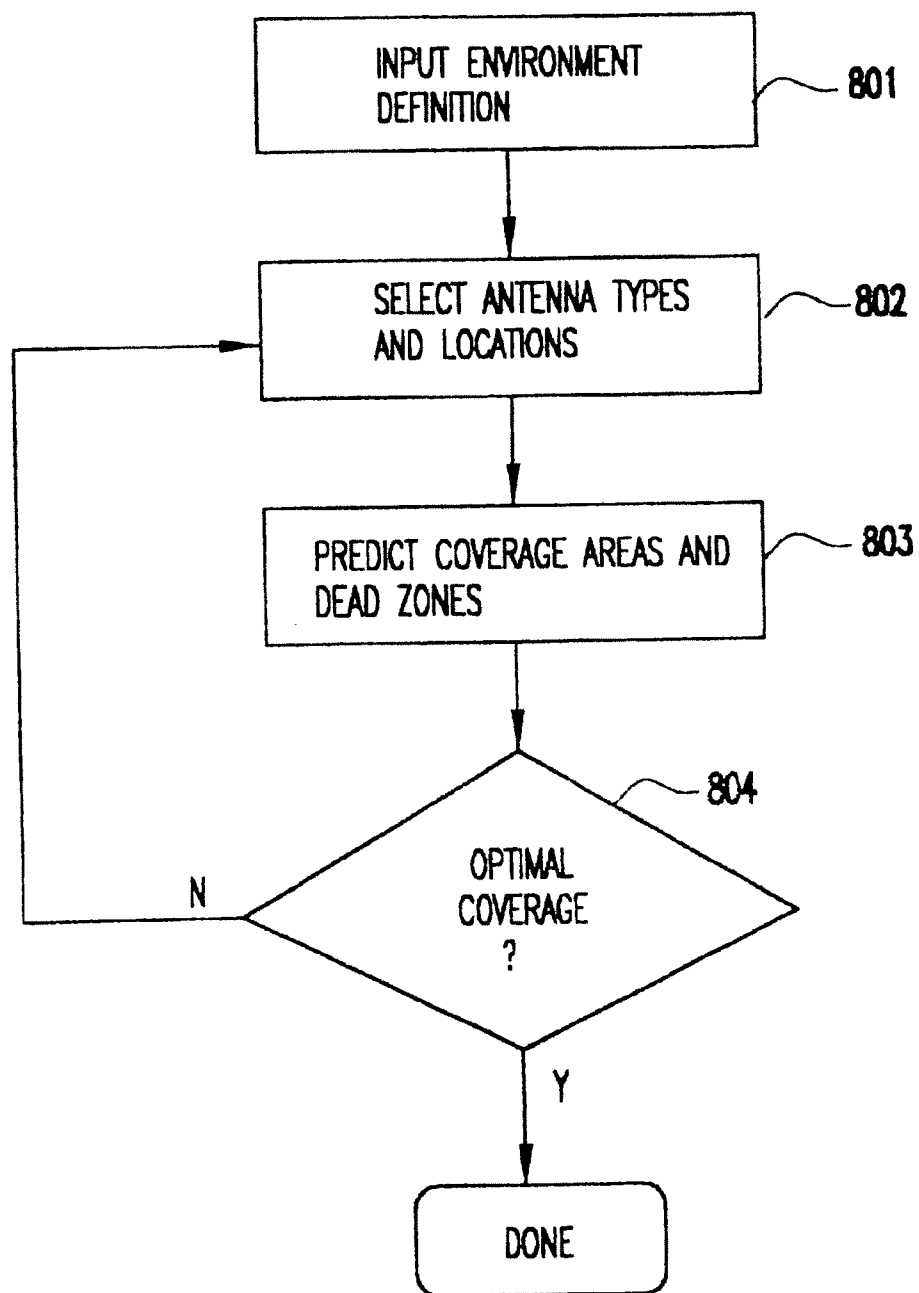


FIG.8

U.S. Patent

Dec. 10, 2002

Sheet 9 of 19

US 6,493,679 B1

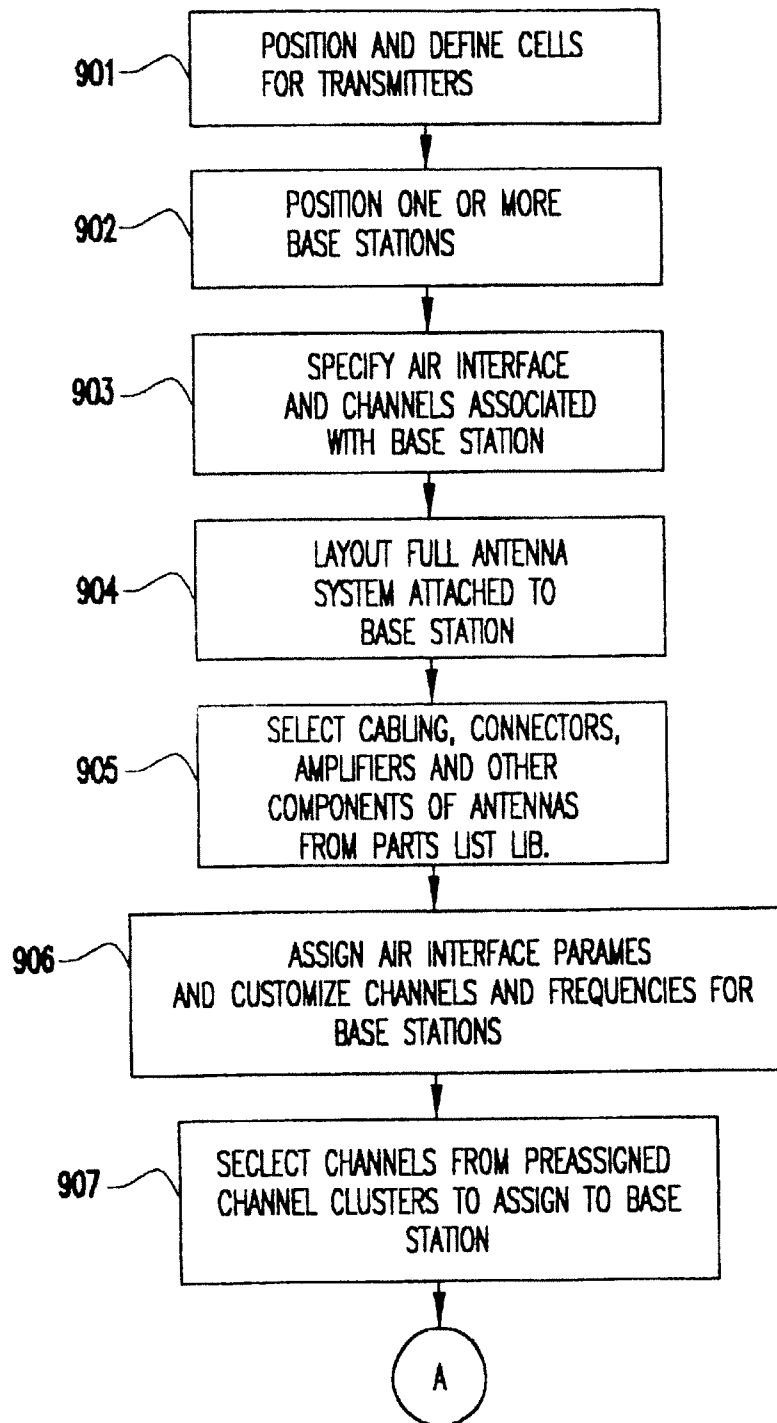


FIG.9A

U.S. Patent

Dec. 10, 2002

Sheet 10 of 19

US 6,493,679 B1

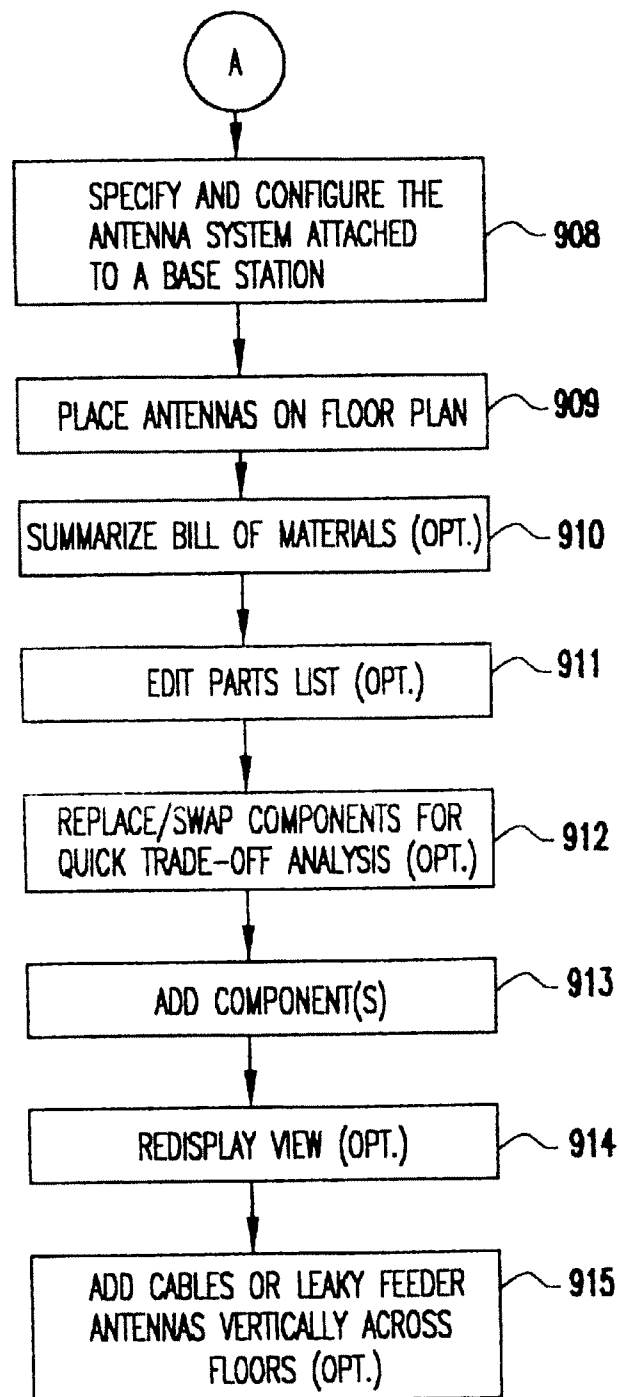


FIG.9B

U.S. Patent

Dec. 10, 2002

Sheet 11 of 19

US 6,493,679 B1

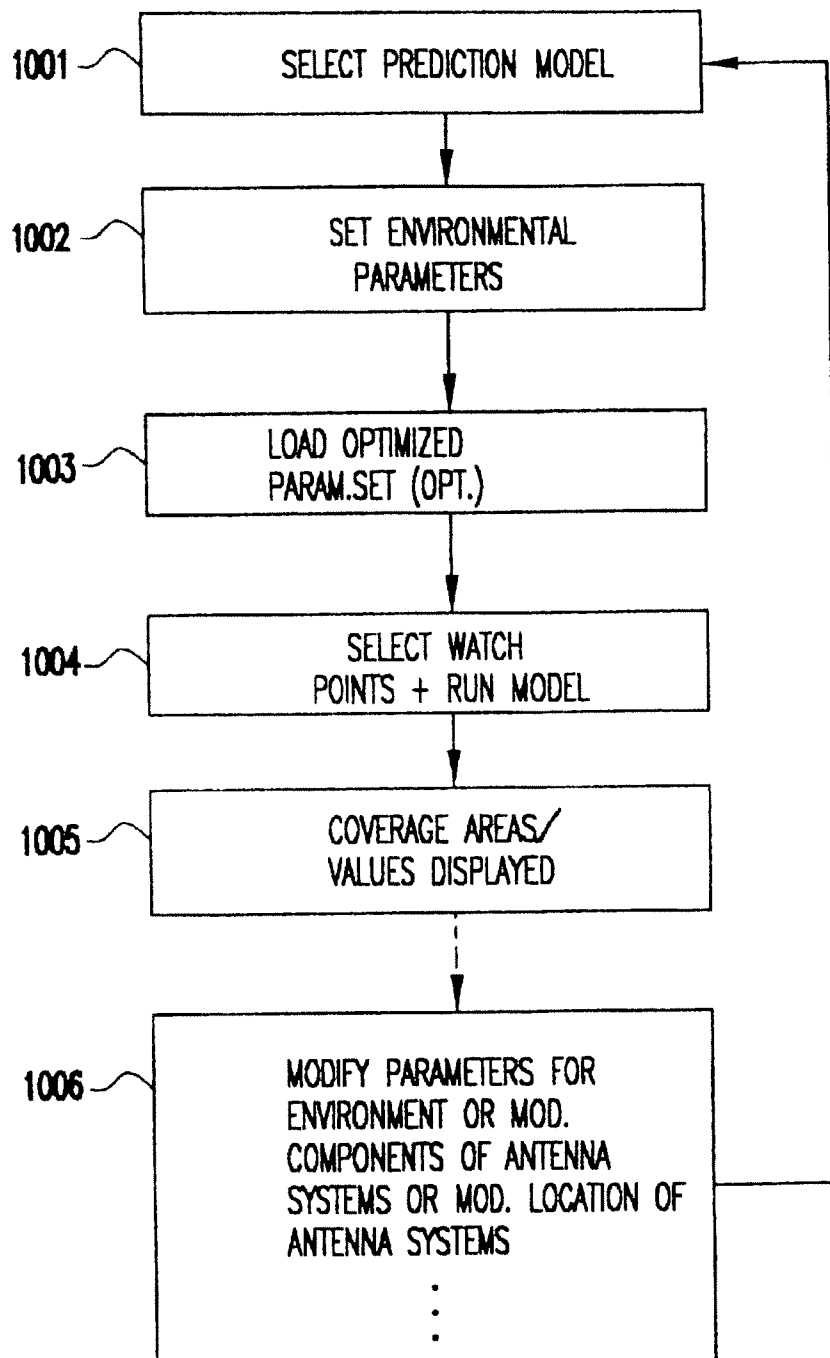


FIG.10

U.S. Patent

Dec. 10, 2002

Sheet 12 of 19

US 6,493,679 B1

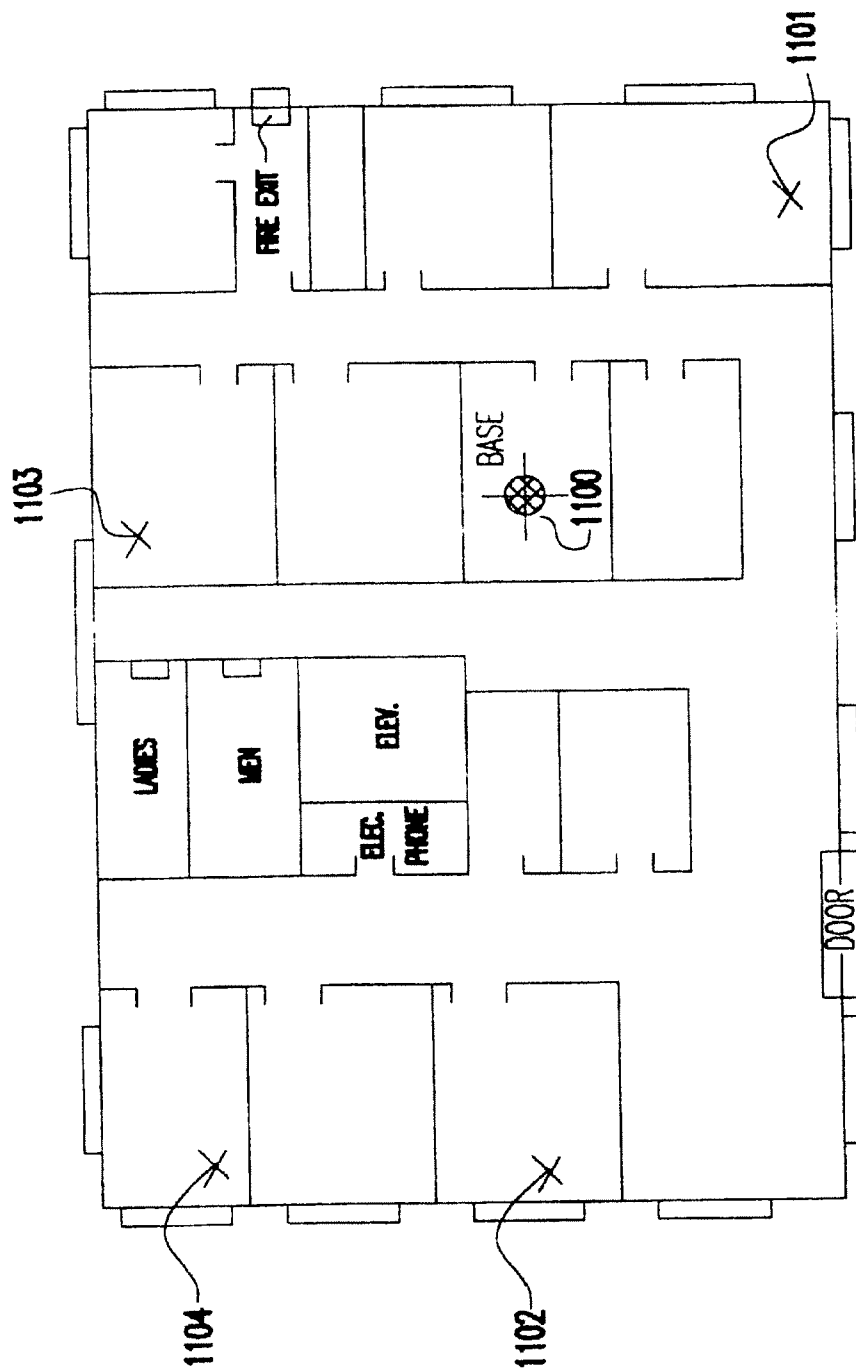


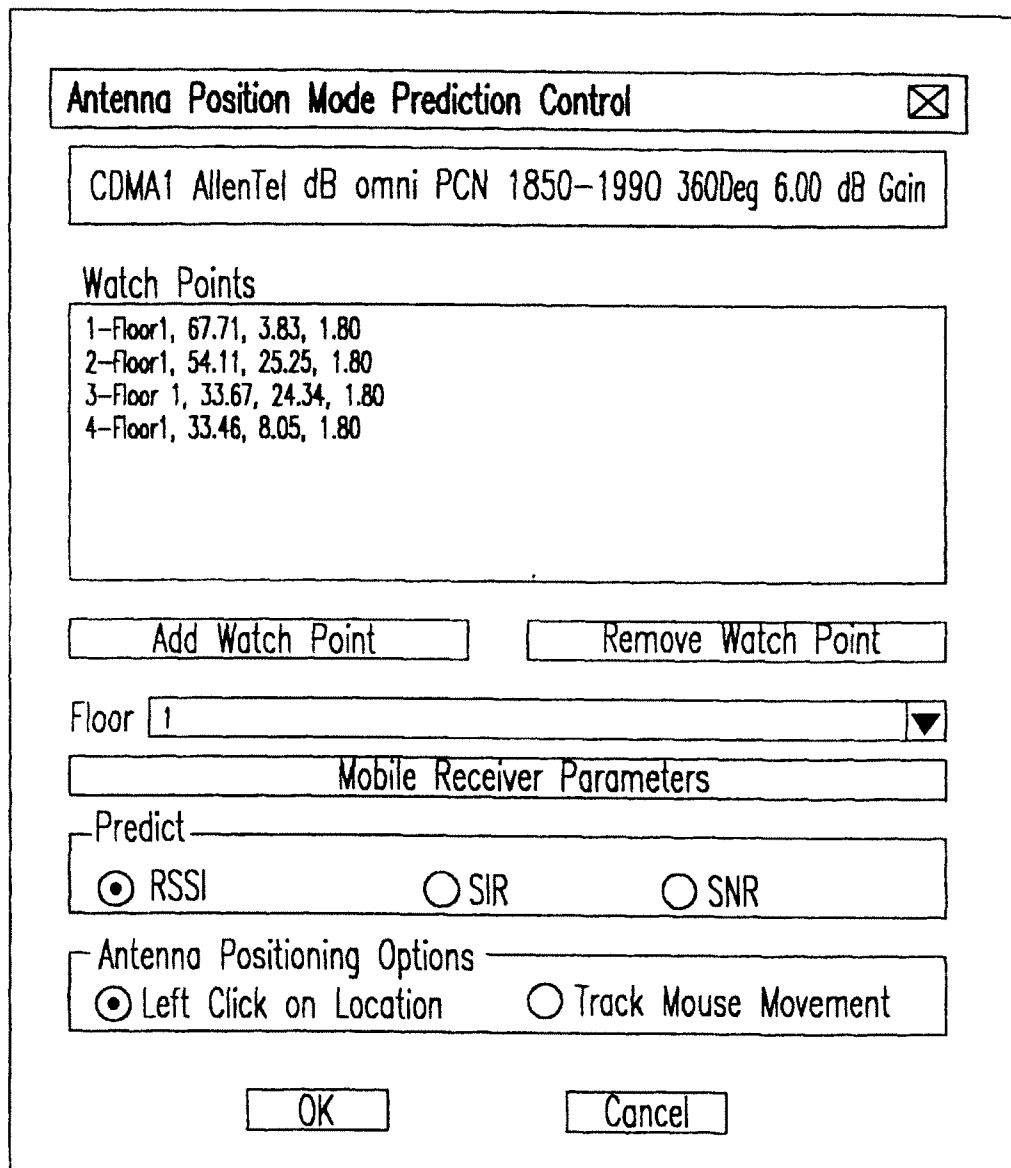
FIG.11

U.S. Patent

Dec. 10, 2002

Sheet 13 of 19

US 6,493,679 B1



Antenna Position Mode Prediction Control ☐

CDMA1 AllenTel dB omni PCN 1850-1990 360Deg 6.00 dB Gain

Watch Points

1-Floor1,	67.71,	3.83,	1.80
2-Floor1,	54.11,	25.25,	1.80
3-Floor 1,	33.67,	24.34,	1.80
4-Floor1,	33.46,	8.05,	1.80

Floor

Mobile Receiver Parameters

Predict

☒ RSSI ☐ SIR ☐ SNR

Antenna Positioning Options

☒ Left Click on Location ☐ Track Mouse Movement

FIG.12

U.S. Patent

Dec. 10, 2002

Sheet 14 of 19

US 6,493,679 B1

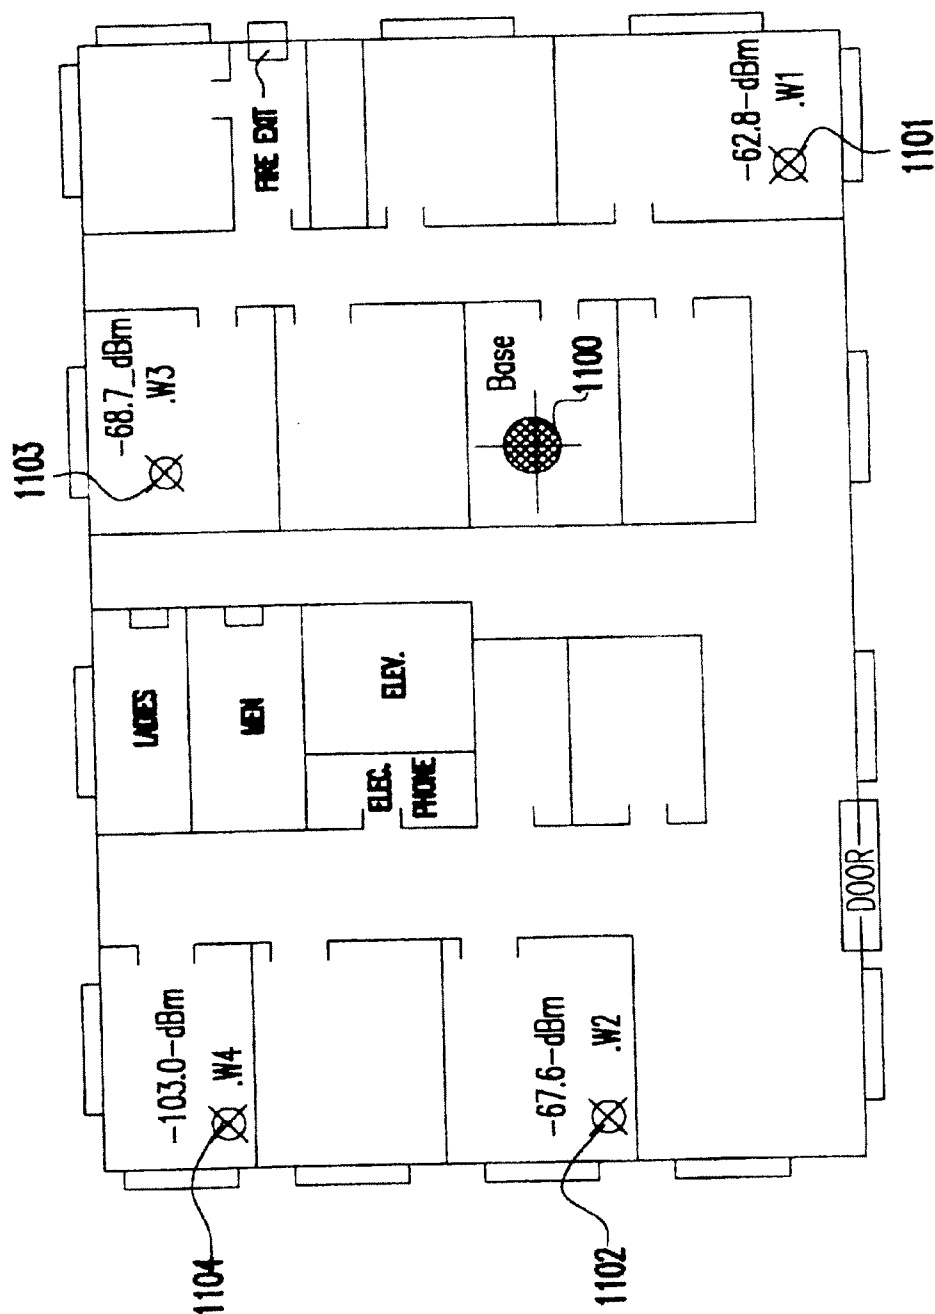


FIG.13

U.S. Patent

Dec. 10, 2002

Sheet 15 of 19

US 6,493,679 B1

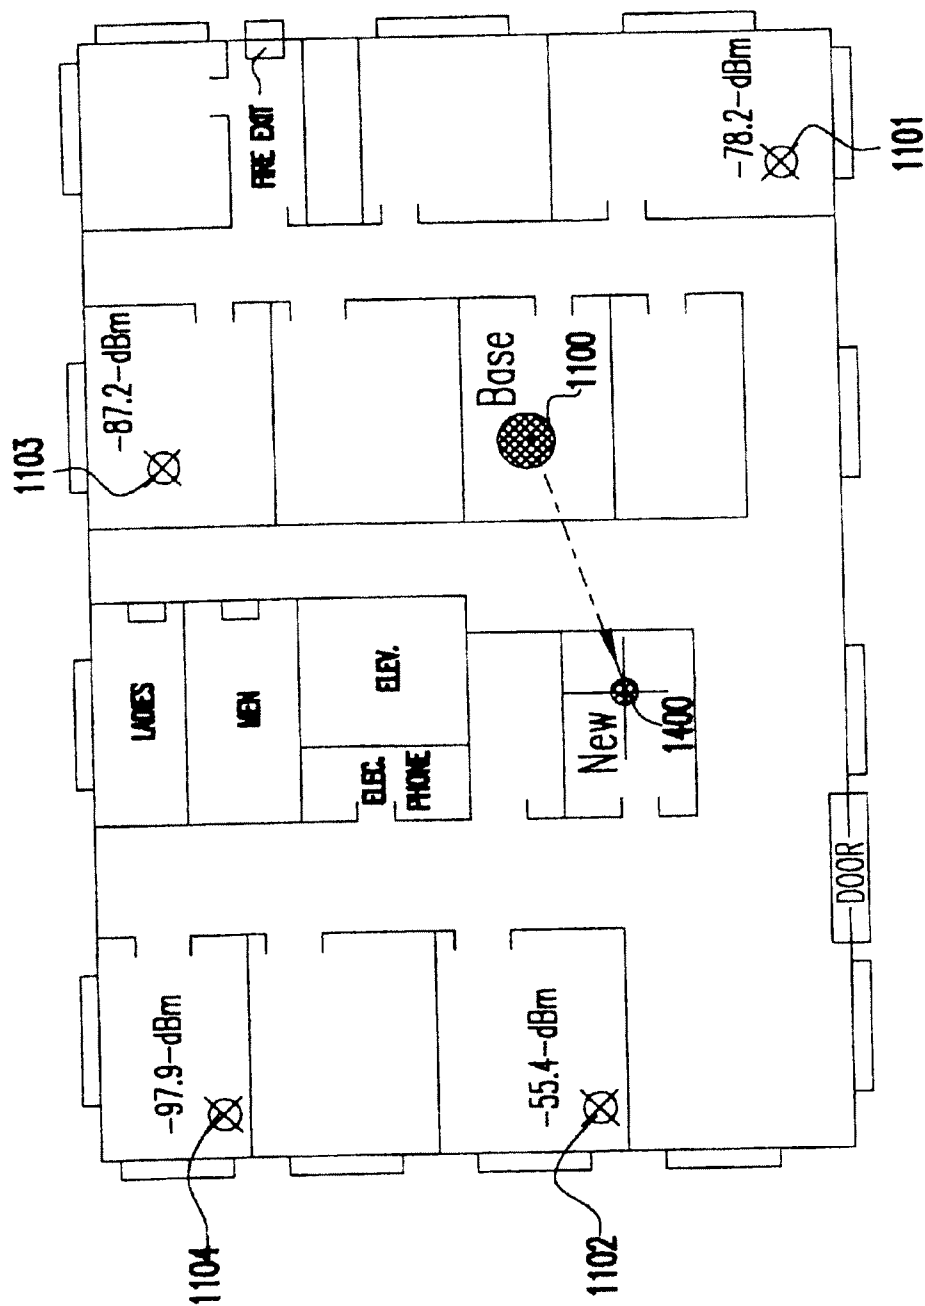


FIG. 14

U.S. Patent

Dec. 10, 2002

Sheet 16 of 19

US 6,493,679 B1

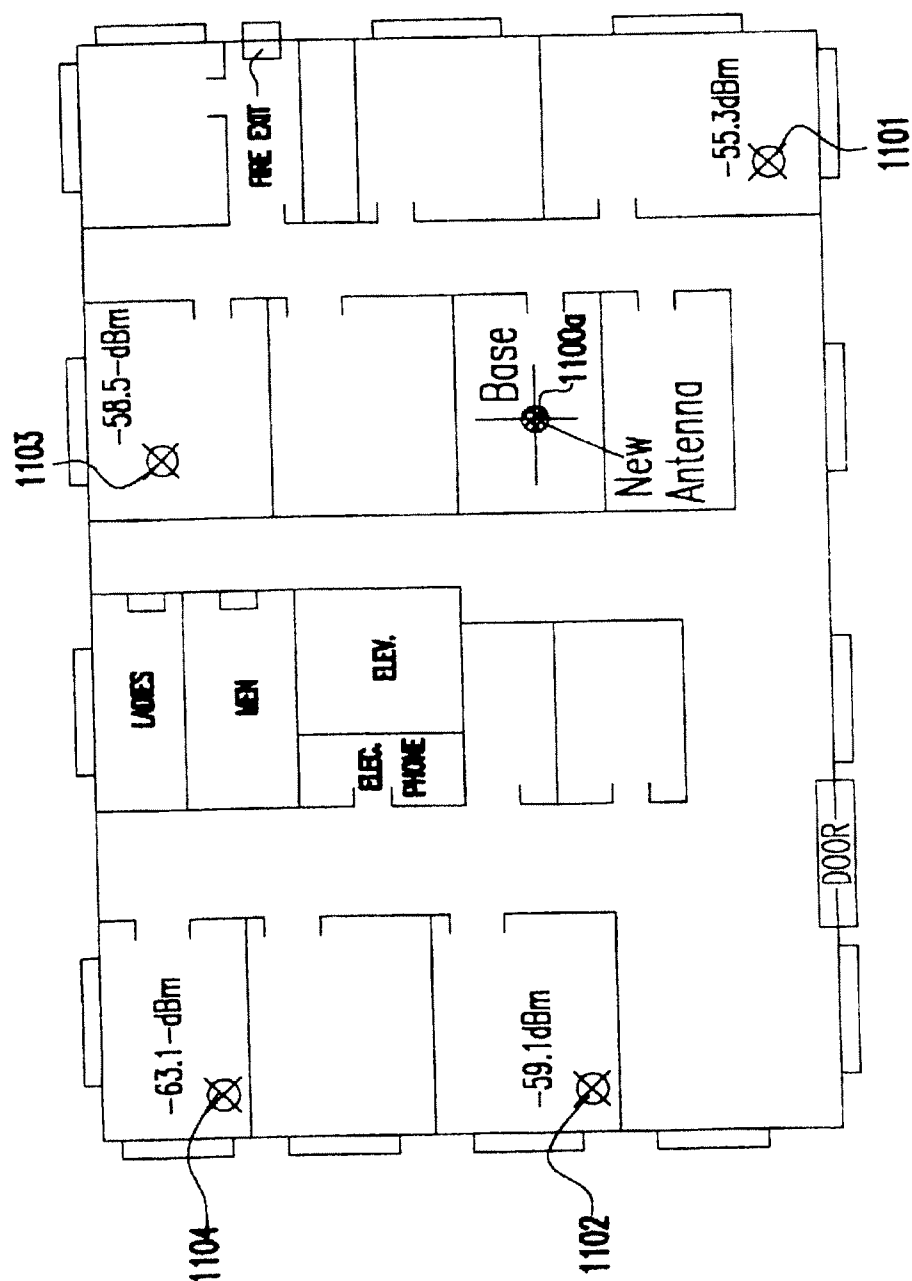


FIG. 15

U.S. Patent

Dec. 10, 2002

Sheet 17 of 19

US 6,493,679 B1

Bill of Materials for Current Drawing

1610 { SUBTOTAL (excluding base station CDMA1): \$0.00

BASE STATION: MACROCELL
DESCRIPTION: CDMA MACROCELL
FLOOR1
POSITION: 84.3, 44.0, 1.8
CHANNEL SET: MACROCELL: IS-95A CDMA Default
SUBCHANNEL SET: Default Channel Set
TXPOWER: 10.00 dBm
RF Bandwidth: 1.25 MHz
RECEIVER NOISE FIGURE: 0.00 dB
CHANNELS ASSIGNED TO BASE STATION
1

1611 { --NAME: AllenTel PCN PANEL 1710-1990 92 Deg 9.00 dB Gain
TYPE: ANTENNA_POINT
MANUFACTURER: Allen Telecom
PART NUMBER: DB972 1850
FREQUENCY: 1710-1990 MHz
PATTERN FILE: 972_185.ant
FLOOR1
POSITION: 84.3, 44.0, 1.8
COST: \$0.00 1612

SUBTOTAL (excluding base station MACROCELL): \$0.00 1613
TOTAL COST(excluding base stations): \$0.00 1614

Save to ASCII File ok

FIG.16

U.S. Patent

Dec. 10, 2002

Sheet 18 of 19

US 6,493,679 B1

Bill of Materials for Current Drawing

1611 {

TYPE: ANTENNA_POINT
 MANUFACTURER: Allen Telecom
 PART NUMBER: DB972 1850
 FREQUENCY: 1710-1990 MHz
 PATTERN FILE: 972_185.ant
 FLOOR1
 POSITION 84.3, 44.0, 1.8
 COST: \$250.00 ~ 1612a

1720 {

NAME: 7/8", 50-ohm Foam Dielectric Coaxial Cable
 TYPE: CABLE
 MANUFACTURER: Andrew
 PART NUMBER: LDF5*
 FREQUENCY: 2000MHz
 LENGTH: 120.41 m (395.06ft)
 LOSS PER 100 m: 6.46 dB
 TOTAL LOSS: 7.78 dB
 POSITION:
 Vertex0: 10.6, 0.8, 1.8
 Vertex1: 1.7, 2.8, 1.8
 Vertex2: 1.7, 31.0, 1.8
 Vertex3: 35.3, 31.0, 1.8
 Vertex4: 35.3, 23.5, 1.8
 Vertex5: 65.4, 23.6, 1.8
 Vertex6: 72.6, 32.0, 1.8
 COST: \$85.00 ~ 1721

SUBTOTAL(excluding base station MACROCELL): \$470.00 ~ 1613a

TOTAL COST(excluding base stations): \$470.00 ~ 1614a

Save to ASCII File ok

FIG. 17

U.S. Patent

Dec. 10, 2002

Sheet 19 of 19

US 6,493,679 B1

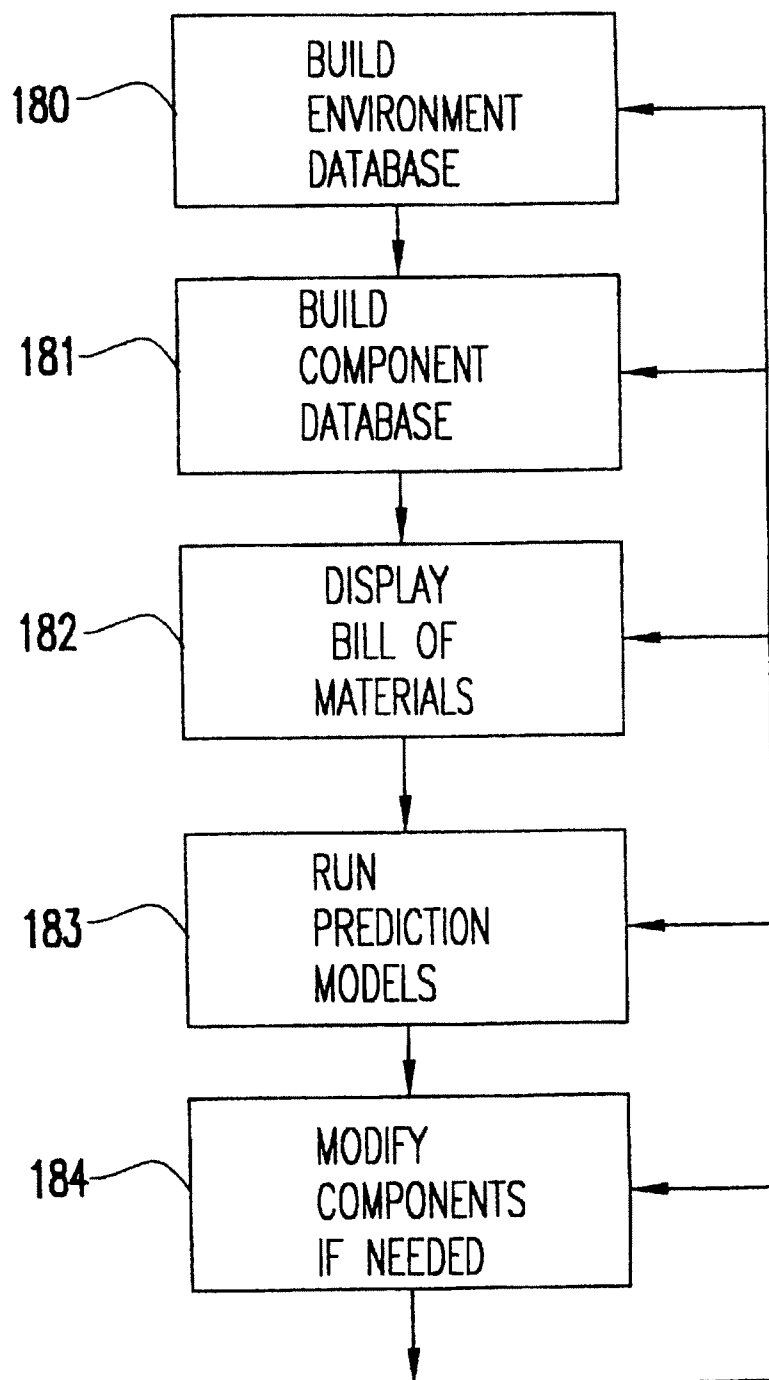


FIG.18

US 6,493,679 B1

1

METHOD AND SYSTEM FOR MANAGING A REAL TIME BILL OF MATERIALS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to concurrently filed applications Ser. No. 09/318,841, entitled "Method And System For a Building Database Manipulator," filed by T. S. Rappaport and R. R. Skidmore and Ser. No. 09/318,840, entitled "Method and System for Automated Optimization of Antenna Positioning in 3-D," filed by T. S. Rappaport and R. R. Skidmore and assigned to a common assignee, the subject matter of which is incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention generally relates to engineering and management systems for the design of wireless communications networks and, more particularly, to a method for managing a real time bill of materials when designing, evaluating or optimizing the performance and/or costs of a wireless communication system using a three-dimensional (3-D) representation of the environment.

2. Background Description

As wireless communications use increases, radio frequency (RF) coverage within buildings and signal penetration into buildings from outside transmitting sources has quickly become an important design issue for wireless engineers who must design and deploy cellular telephone systems, paging systems, or new wireless systems and technologies such as personal communication networks or wireless local area networks. Designers are frequently requested to determine if a radio transceiver location, or base station cell site can provide reliable service throughout an entire city, an office, building, arena or campus. A common problem for wireless systems is inadequate coverage, or a "dead zone," in a specific location, such as a conference room. It is now understood that an indoor wireless PBX (private branch exchange) system or wireless local area network (WLAN) can be rendered useless by interference from nearby, similar systems. The costs of in-building and microcell devices which provide wireless coverage within a 2 kilometer radius are diminishing, and the workload for RF engineers and technicians to install these on-premises systems is increasing sharply. Rapid engineering design and deployment methods for microcell and in-building wireless systems are vital for cost-efficient build-out.

Analyzing radio signal coverage penetration and interference is of critical importance for a number of reasons. A design engineer must determine if an existing outdoor large scale wireless system, or macrocell, will provide sufficient coverage throughout a building, or group of buildings (i.e., a campus). Alternatively, wireless engineers must determine whether local area coverage will be adequately supplemented by other existing macrocells, or whether indoor wireless transceivers, or picocells, must be added. The placement of these cells is critical from both a cost and performance standpoint. If an indoor wireless system is being planned that interferes with signals from an outdoor macrocell, the design engineer must predict how much interference can be expected and where it will manifest itself within the building, or group of buildings. Also, providing a wireless system that minimizes equipment infrastructure cost as well as installation cost is of significant economic importance. As in-building and microcell wireless systems proliferate, these issues must be resolved quickly, easily, and inexpensively, in a systematic and repeatable manner.

2

There are many computer aided design (CAD) products on the market that can be used to design the environment used in one's place of business or campus. WISE from Lucent Technology, Inc., SignalPro from EDX, PLANet by Mobile Systems International, Inc., and TEMS and TEMS Light from Ericsson are examples of wireless CAD products. In practice, however, a pre-existing building or campus is designed only on paper and a database of parameters defining the environment does not readily exist. It has been difficult, if not generally impossible, to gather this disparate information and manipulate the data for the purposes of planning and implementation of indoor and outdoor RF wireless communication systems, and each new environment requires tedious manual data formatting in order to run with computer generated wireless prediction models. Recent research efforts by AT&T Laboratories, Brooklyn Polytechnic, and Virginia Tech, are described in papers and technical reports entitled "Radio Propagation Measurements and Prediction Using Three-dimensional Ray Tracing in Urban Environments at 908 MHz and 1.9 GHz," (IEEE Transactions on Vehicular Technology, VOL. 48, No. 3, May 1999), by S. Kim, B. J. Guarino, Jr., T. M. Willis III, V. Erceg, S. J. Fortune, R. A. Valenzuela, L. W. Thomas, J. Ling, and J. D. Moore, (hereinafter "Radio Propagation"); "Achievable Accuracy of Site-Specific Path-Loss Predictions in Residential Environments," (IEEE Transactions on Vehicular Technology, VOL. 48, No. 3, May 1999), by L. Piazzzi and H. L. Bertoni; "Measurements and Models for Radio Path Loss and Penetration Loss In and Around Homes and Trees at 5.85 Ghz," (IEEE Transactions on Communications, Vol. 46, No. 11, November 1998), by G. Durgin, T. S. Rappaport, and H. Xu; "Radio Propagation Prediction Techniques and Computer-Aided Channel Modeling for Embedded Wireless Microsystems," ARPA Annual Report, MPRG Technical Report MPRG-TR-94-12, July 1994, 14 pp., Virginia Tech, Blacksburg, by T. S. Rappaport, M. P. Koushik, J. C. Liberti, C. Pendyala, and T. P. Subramanian; "Radio Propagation Prediction Techniques and Computer-Aided Channel Modeling for Embedded Wireless Microsystems," MPRG Technical Report MPRG-TR-95-08, July 1995, 13 pp., Virginia Tech, Blacksburg, by T. S. Rappaport, M. P. Koushik, C. Carter, and M. Ahmed; "Use of Topographic Maps with Building Information to Determine Antenna Placements and GPS Satellite Coverage for Radio Detection & Tracking in Urban Environments," MPRG Technical Report MPRG-TR-95-14, Sep. 15, 1995, 27 pp., Virginia Tech, Blacksburg, by T. S. Rappaport, M. P. Koushik, M. Ahmed, C. Carter, B. Newhall, and N. Zhang; "Use of Topographic Maps with Building Information to Determine Antenna Placement for Radio Detection and Tracking in Urban Environments," MPRG Technical Report MPRG-TR-95-19, November 1995, 184 pp., Virginia Tech, Blacksburg, by M. Ahmed, K. Blankenship, C. Carter, P. Koushik, W. Newhall, R. Skidmore, N. Zhang and T. S. Rappaport; "A Comprehensive In-Building and Microcellular Wireless Communications System Design Tool," MPRG-TR-97-13, June 1997, 122 pp., Virginia Tech, Blacksburg, by R. R. Skidmore and T. S. Rappaport; "Predicted Path Loss for Rosslyn, V. A.," MPRG-TR-94-20, Dec. 9, 1994, 19 pp., Virginia Tech, Blacksburg, by S. Sandhu, P. Koushik, and T. S. Rappaport; "Predicted Path Loss for Rosslyn, Va., Second set of predictions for ORD Project on Site Specific Propagation Prediction" MPRG-TR-95-03, Mar. 5, 1995, 51 pp., Virginia Tech, Blacksburg, by S. Sandhu, P. Koushik, and T. S. Rappaport. These papers and technical reports are illustrative of the state of the art in site-specific propagation modeling and show the difficulty in obtaining databases for

US 6,493,679 B1

3

city environments, such as Rosslyn, Va. While the above papers describe a research comparison of measured vs. predicted signal coverage, the works do not demonstrate a systematic, repeatable and fast methodology for creating an environmental database, nor do they report a method for analyzing system performance or visualizing and placing various wireless equipment components that are required to provide signals in the deployment of a wireless system in that environment.

While there are many methods available for designing wireless networks that provide adequate coverage, there is no easy method to ensure that the system will be cost effective. For instance, even though the coverage may be more than adequate, given the chosen wireless infrastructure components, the total cost of the system could be prohibitive.

SUMMARY OF THE INVENTION

It is an object of the invention to provide a rapid and automated method for generating a bill of materials and cost information in real time, as components for a desired wireless communication system are specified and/or replaced by substitute components, while continuously predicting wireless system performance. This automatic method for comparing the cost and performance of competing products or competing design methodologies, in real time, offers a significant value for wireless engineers and provides a marked improvement over present day techniques.

According to the invention, a design engineer builds a model of the desired wireless communications system and specifies each component necessary to provide sufficient or optimal system performance. A parts list is maintained, in real time, that contains a definition of each system component and its associated performance and cost parameters. Using this method, the user is able to rapidly change the physical location of components within the wireless system in order to investigate alternative designs which may use different components, such as antennas, cables; or use different RF distribution methods and/or various types of coaxial or optical splitter systems, etc. Cost parameters include both component costs and installation costs. As the system is changed through a series of "whatif" scenarios, components are replaced with substitute components cable lengths are modified, antenna systems and base stations are re-positioned to alternate locations, etc.

Each time a component is added to or deleted from the system model, the bill of materials is automatically updated and component costs, total costs, and altered system performance specifications are immediately available to the design engineer. The designer may choose to swap components for less expensive components. The performance characteristics of the system are automatically updated as cost choices are made to enable the designer to assess the changes in performance and cost at the same time.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

FIG. 1 shows an example of a simplified layout of a floor plan of a building;

FIG. 2 shows effective penetration of Radio Frequency (RF) transmission into a building from a macrocell;

FIG. 3 shows a simplified layout of a floor plan of a building including an outdoor macrocell and an indoor base station;

4

FIG. 4 shows the layout of FIG. 3, but with a revised base station designed to eliminate interference;

FIG. 5 is a flow diagram of a general method used to design a wireless communication network;

FIG. 6 is a flow diagram of a method used to generate estimates based on field measurements;

FIG. 7 is a flow diagram of a method used to match best propagation parameters with measured data;

FIG. 8 is a flow diagram of a method for prediction;

FIGS. 9A and 9B together make up a detailed flow diagram of a method to generate a design of a wireless network and determine its adequacy;

FIG. 10 is a flow diagram showing a method for using watch points during antenna repositioning or modification;

FIG. 11 shows a simplified layout of a floor plan of a building with a base station and watch points selected;

FIG. 12 shows a dialog box displaying the locations of the selected watch points and choices for display information;

FIG. 13 shows a simplified layout of a floor plan of a building with a base station and initial RSSI values for the selected watch points;

FIG. 14 shows a simplified layout of a floor plan of a building with a repositioned base station and changed RSSI values for the selected watch points;

FIG. 15 shows a simplified layout of a floor plan of a building with a re-engineered base station and changed RSSI values for the selected watch points;

FIG. 16 shows a bill of materials summary for a drawing, according to the preferred embodiment of the invention;

FIG. 17 shows a bill of materials summary for a drawing after costs have been added to a database, according to the preferred embodiment of the invention; and

FIG. 18 is a flow diagram showing the general method of the present invention.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT OF THE INVENTION

Design of Wireless Communication Systems

Using the present method, it is now possible to assess the RF environment in a systematic, organized fashion by quickly viewing signal strength, or interference levels, or other wireless system performance measures. The current embodiment is designed specifically for use with the SitePlanner™ suite of products available from Wireless Valley Communications, Inc. of Blacksburg, Va. However, it will be apparent to one skilled in the art that the method could be practiced with other products either now known or to be developed in the future. (SitePlanner is a trademark of Wireless Valley Communications, Inc.)

Referring now to FIG. 1, there is shown a two-dimensional (2-D) simplified example of a layout of a building floor plan. The method uses 3-D computer aided design (CAD) renditions of a building, or a collection of buildings and/or surrounding terrain and foliage. However, for simplicity of illustration a 2-D figure is used. The various physical objects within the environment such as external walls 101, internal walls 102 and floors 103 are assigned appropriate physical, electrical, and aesthetic values. For example, outside walls 101 may be given a 10 dB attenuation loss, signals passing through interior walls 102 may be assigned 3 dB attenuation loss, and windows 104 may show a 2 dB RF penetration loss. In addition to attenuation, the obstructions 101, 102 and 103 are assigned other properties including reflectivity and surface roughness.

US 6,493,679 B1

5

Estimated partition electrical properties loss values can be extracted from extensive Propagation measurements already published, which are deduced from field experience, or the partition losses of a particular object can be measured directly and optimized instantly using the present invention combined with those methods described in the copending application Ser. No. 09/221,985, entitled "System for Creating a Computer Model and Measurement Database of a Wireless Communication Network" filed by T. S. Rappaport and R. R. Skidmore. Once the appropriate physical and electrical parameters are specified, any desired number of hardware components of RF sources can be placed in the 3-D building database, and received signal strengths (RSSI), network throughput, bit or frame error rate, or carrier-to-interference (C/I) ratios can be plotted directly onto the CAD drawing. The 3-D environment database could be built by a number of methods, the preferred method being disclosed in the concurrently filed, copending application Ser. No. 09/318,841, (Docket 256015AA). Traffic capacity analysis, frequency planning, co-channel interference analysis can be performed in the invention along with RF coverage. Other system performance metrics may be easily incorporated by one skilled in the art through well known equations.

FIG. 2 depicts effective RF penetration into the building from the distant macrocell using a close-in virtual macrocell transmitting into the lossless distributed antenna.

Referring to FIG. 2, there are several windows 104, and even a large glass foyer 105, on the north wall of the building, so RF penetration into this part of the building is quite good, as shown by contour lines 108 and 109 for 0 dB and -30 dB, respectively. Even so, interior walls 102 cause signal levels in some areas to drop below a minimum useable signal strength of about -90 dBm, especially in some of the southern rooms, as shown by contour line 110. Consequently, macrocell coverage there will probably be poor.

Other outdoor macrocells can be modeled in the same way, and their signal strength contours plotted, to determine if hand-offs can compensate for the inadequacies of the macrocell north of the building. If not, then indoor picocells (and their distributed feed systems, antennas, and antenna patterns) can be easily added if necessary, and their performance checked using the method, to complement coverage provided by the macrocells.

The mathematical propagation models used to predict and optimize antenna positioning in a desired environment may include a number of predictive techniques models, such as those described in the previously cited and following technical reports and papers: "Interactive Coverage Region and System Design Simulation for Wireless Communication Systems in Multi-floored Indoor Environments, SMT Plus," IEEE ICUPC '96 Proceedings, by R. R. Skidmore, T. S. Rappaport, and L. Abbott which is hereby incorporated by reference. Some simple models are also briefly described in "SitePlanner 3.16 for Windows 95/98/NT User's Manual" (Wireless Valley Communications, Inc. 1999), hereby incorporated by reference. It would be apparent to one skilled in the art how to apply other system performance models to this method.

Interference, instead of radio signal strength, is the dominant performance-limiting factor in many situations due to increased wireless communications use. Modeling interference from any source to an established or contemplated wireless system is straightforward using the method. Suppose, for example, that an indoor wireless communica-

6

tion system is assigned a frequency set identical to that of an outdoor wireless system. Although the indoor system may provide sufficient RSSI throughout its coverage area, interference from the outside system may still render the indoor wireless system ineffectual in certain parts of the building.

Caution must be used, however, when modeling and analyzing interference, since the detrimental effect may also depend upon technologies and/or signal processing technologies, not just signal power levels. For example, a geographic area could have similar narrowband and/or wideband in the 800 MHZ cellular bands, for instance with Advanced Mobile Phone System (AMPS) and Code Division Multiple Access (CDMA) systems, but users using either technology may be able to coexist if their respective demodulation processes reject interference provided by the undesired system. The current embodiment of this invention allows the user to select the air interface/technology being used by the wireless system being designed and automatically adjusts the prediction of interference accordingly.

FIG. 3 shows another rendition of the office building example, but an indoor wireless system 107 has been added. In this example, 800 MHZ AMPS technology is assigned to both transmitters 106 and 107. Differing wireless standards and technologies such as CDMA and Global System Mobile (GSM) could have been selected as well. The present invention uses a database to represent the exact physical air interface standards of a wide range of technologies and may be easily edited for future interface standards. As new technologies are developed, one skilled in the art could easily modify this invention to include the new technologies.

The outdoor wireless system 106 is now interfering with the indoor network, and the effect is checked by plotting C/I contours 111 and 112 at 0 dB and -30 dB, respectively, for the outdoor system and also plotting C/I contours 113 and 114 at 0 dB and -30 dB for the indoor system. The 0 dB contour 114 shows where the desired and interfering signal levels are equal, so the interfering outdoor system's signal predominates in areas outside this contour. It is obvious that the indoor network is rendered useless throughout many parts of the building. There are a number of possible solutions that may be analyzed by a designer using the present invention.

One solution is to change the indoor system's antenna location or increase the transmitted power, add more nodes, or select a different frequency set. These changes may be made with the simple click of a mouse in the method of the invention, so that new channel sets, antenna locations, or alternative antenna systems (such as in-building distributed systems, directional antennas, or leaky feeders) may be evaluated quickly, thereby eliminating guesswork and/or costly on-site experimentation with actual hardware. Instead of displaying contours of coverage or interference, the present invention also allows the user to specify fixed or moveable watch points that indicate or display predicted performance in extremely rapid fashion at specific points in the environment.

For example, FIG. 4 illustrates how the same indoor wireless system of FIG. 3 can provide adequate C/I protection when connected to a distributed indoor antenna system consisting of a two-way splitter 401 (3 dB loss+ insertion loss) and two 40 foot cable runs 402 to popular commercial indoor omnidirectional antennas 403. A look at the new 0 dB contour lines 111 and 215, and -30 dB contour lines 112a and 216 show that the coverage inside the building is now adequate; the outdoor system 106 no longer causes significant interference in most parts of the building. Watch points

US 6,493,679 B1

7

allow a user to instantly determine spot coverage or other system performance without having to wait for the computation and display of contour plots.

The method allows any type of distributed antenna system to be modeled within seconds, while continuously monitoring and analyzing the component and installation cost and resulting link budget, as disclosed below, enabling "what-if" designs to be carried out on the fly with minimum guess work and wasted time.

In the present embodiment of the invention, the designer identifies locations in the 3-D environmental database where certain levels of wireless system performance are desirable or critical. These locations, termed "watch points", are points in three-dimensional space which the designer identifies by visually pointing and/or clicking with a mouse or other input device at the desired location in the 3-D environmental database. Any number of such watch points may be placed throughout the 3-D environment at any location. Watch points may be designated prior to performing a performance prediction on a given wireless communication system, or may be dynamically created by the user at any time during the course of a wireless system performance calculation using the same point and click technique described above.

Watch points provide graphical and/or textual feedback to a designer regarding the wireless system performance throughout the environment. Depending on the type of visual feedback desired by the designer, watch points may take the form of one or more of the following:

- A computed number displayed as text that represents received signal strength (RSSI), signal-to-interference ratio (SIR), signal-to-noise ratio (SNR), frame error rate (FER), bit error rate (BER), or other wireless system performance metrics;

- A small region of solid color whose shade and/or tint varies relative to some computed wireless system performance metric;

- Colored lines linking the watch point location with the location one or more antennas in the wireless communication system, where the color, thickness, and/or other physical aspect of the connecting line varies relative to some computed wireless system performance metric and dependent upon whether the forward or reverse wireless system channel is being analyzed;

- Other form designated by the designer; or

- Any combination of the above.

In all cases, the graphical and/or textual representation of each watch point is updated in real-time as a result of the instantaneous computation of the wireless system performance metrics, which are linked to the 3-D environmental database, and initiated due to dynamic changes being made to the wireless system configuration and/or watch point position itself by the user. For example, if the user repositions an antenna using the mouse or other input device, the effect of doing so on the overall wireless system performance is computed and the results are displayed via changes in the appearance of watch points. In addition, numerical values predicted at the watch points are displayed in summary in a dialog window and written to a text file for later analysis. This process is described in greater detail in the following sections.

The preferred embodiment of the invention utilizes a 3-D environmental database containing information relevant to the prediction of wireless communication system performance. This information includes but is not limited to the location, and the physical and electromagnetic properties of

8

obstructions within the 3-D environment, where an obstruction could be any physical object or landscape feature within the environment (e.g., walls, doors, windows, buildings, trees, terrain features, etc.), as well as the position and physical and electrical properties of communications hardware to be used or simulated in the environment.

The designer identifies the location and type of all wireless communication system equipment within the 3-D environmental database. This point-and-click process involves the designer selecting the desired component from a computer parts database and then visually positioning, orienting, and interconnecting various hardware components within the 3-D environmental database to form complete wireless communication systems. The preferred embodiment of the computer parts database is more fully described below. The resulting interconnected network of RF hardware components (commonly known as a wireless distributed antenna) is preferably assembled using either a drag and drop technique or a pick and place and is graphically displayed overlaid upon the 3-D environmental database, and utilizes electromechanical information available for each component via the parts list library in order to fully describe the physical operating characteristics of the wireless communication system (e.g., the system noise figure, antenna radiation characteristics, frequencies, etc.). This information is directly utilized during the prediction of wireless system performance metrics and is discussed later.

The present invention represents a dramatic improvement over prior art by providing the design engineer with instant feedback on wireless system performance metrics as the user alters the physical location transmitter, receivers, and other components, or otherwise modifies the antenna system. The current embodiment utilizes the concept of watch points to implement this. Multiple methods of display and a wide range of settings are available for the designer to use in optimizing antenna placement based upon wireless system performance values displayed at each watch point. One skilled in the art could see how watch points as they are herein described could apply to different implementations as well. Descriptions of the different techniques implemented in the current invention are provided in the following sections.

One form of the method allows the designer to dynamically alter the position, orientation, and/or type of any hardware component utilized within a wireless communication system modeled in a 3-D environmental database. Using this technique, the designer may identify watch points representing critical areas of the 3-D environment that require a certain level of wireless system performance. Such areas could include the office of the Chief Executive Officer (CEO) of a company, a conference room, a city park, or the office of a surgeon on call. Next the designer selects the component of interest within the wireless system. In the present invention, this would be the selection of an antenna or leaky feeder antenna, for example, but one skilled in the art could see that this could be any physical antenna system component. Once the desired hardware component is selected, the designer may begin making changes to the state of the component. For example, by moving the mouse or other input device cursor, the user could effectively relocate the selected component to another position in the 3-D environmental database. This involves the user visually moving the mouse cursor, in real-time, such that the cursor resides in another portions of the 3-D database. The present invention recalculates wireless system performance based upon RSSI, SIR, SNR, FER, BER, or other metric, incorporating the user's desired change in the position of the selected component.

US 6,493,679 B1

9

The calculations combine the electromechanical properties of each component in the wireless communication system (e.g., noise figure, attenuation loss or amplification, antenna radiation pattern, etc.), the electromagnetic properties of the 3-D environmental database, and radio wave propagation techniques (detailed later) to provide an estimate of the wireless system performance. Calculations are performed at each watch point the user has identified, and the graphical display of the watch point is updated to reflect the result of the calculations.

As the user moves the mouse cursor and effectively repositions the selected component, the overall performance of the wireless communication system may be altered. For example, if the selected component is an antenna, repositioning the antenna changes the origination point of radio wave signals being broadcast from the antenna, and can thus dramatically change the reception of adequate RF signal throughout the environment. Because the graphical display of the watch points is updated in real-time as the selected component is repositioned, the designer is provided instant feedback on the revised wireless system performance, and can make design decisions based upon the viability of multiple proposed locations and/or wireless system configurations rapidly.

In addition to the functionality described above, the designer is free to add additional watch points in any location within the 3-D environmental database at any time during a wireless system performance prediction. In the current implementation, the designer clicks with the mouse or other input device on the desired location in the 3-D environmental database to create a new watch point at the selected location that is then updated throughout the remainder of the performance prediction.

In a similar fashion, the preferred embodiment enables a designer to reorient a selected antenna in real-time with respect to any coordinate axis while the graphical display of all drawing watch points is updated to reflect the revised wireless system performance metrics as a result of the new antenna orientation.

In a similar fashion, a designer may replace an existing hardware component in the wireless communication system with any component available from the parts list library. In doing so, the changes to the wireless communication system performance as a result of the replacement is reflected in the graphical display of the watch points.

In a similar fashion, a designer may selectively include or exclude any subset of components within the wireless communication system while selecting components to involve in the wireless system performance calculation. For example, a designer could consider the effect of repositioning a single antenna, or could consider the combined, composite effect on the watch points as individual antennas are repositioned within a wireless system network consisting of additional, fixed antenna placements.

In a similar fashion, the designer may choose to allow watch points to be mobile. That is, instead of positioning a watch point and having the graphical display of the watch point reflect the changing wireless system performance metric, the designer could instead identify a watch point whose position is mobile but whose graphical display remains constant. In this scenario, the position of the watch point fluctuates along a linear path traced between itself and the current location of the mouse cursor until a position within the 3-D database is found at which the desired level of wireless system performance metric is maintained. For example, the designer may create a watch point to maintain a constant graphical display synonymous with -65 dBm

10

RSSI. As the user repositions, reorients, or otherwise alters components within the wireless communication system, the watch point alters its position within the 3-D environmental database until a position is found at which a calculated value of -65 dBm RSSI is determined.

In addition to enabling a designer to reposition, reorient, and/or replace wireless system components in real-time while visualizing the impact of such changes at selected watch points within the 3-D database, the user may choose to maintain the current configuration of the wireless communication system and instead create a single, mobile watch point. The watch point thus created is dynamically repositioned within the 3-D environmental database in real-time by the user by simply repositioning the mouse cursor. Positioning the mouse cursor at a given location within the 3-D environmental database is equivalent to repositioning the watch point to match that location. In the present invention, this technique is used to allow the mobile watch point to represent a mobile user in the 3-D environmental database. As in the previous scenarios, the graphical display of the watch point is updated in real-time to reflect predicted wireless system performance metrics at the watch point position. The designer is free to select individual subsets of wireless system components to involve in the calculations of wireless system performance. Thus the graphical display of the watch point may reflect the performance metrics specific to individual wireless system components or the composite performance metrics due to the combined effect of multiple selected components. For example, the radiating power of multiple antennas can be combined into a single measure of received signal strength.

The two primary uses of the single mobile watch point technique involve the analysis of the forward link (or down link) and reverse link (or up link) of the wireless system. The forward link of a wireless communication system involves the flow of radio signals from the fixed wireless system to the mobile user, while the reverse link of a wireless communication system involves the flow of radio signals from the mobile user to the fixed wireless system. In the present embodiment, line segments are drawn between the mobile watch point (which is also the mouse cursor) to each antenna the designer has included in the wireless system performance prediction. In addition, the individual or subsets of antennas identified as having the best wireless system performance characteristics are differentiated from the other antennas by altering the color and/or other physical appearance of the connector lines between the antennas and the watch point. As the designer then repositions the mouse cursor, the selected location for the watch point in the 3-D database, and therefore the effective location of the mobile user, is adjusted to match that of the mouse cursor. The wireless system performance metrics are recalculated at the watch point location for the antenna components selected by the designer, and the graphical display of the watch point and all connector lines is updated accordingly.

Another improvement over the prior art is the ability to dynamically model the repositioning of leaky feeder antennas and visualize the effects on wireless system performance. A leaky feeder antenna can be thought of as a cable with many holes regularly spaced along its length. Such a cable would experience a signal loss or emanation at every hole and would thus radiate RF energy along the entire cable length. Leaky feeder antenna, or lossy coaxial cable as it is sometimes referred, can be thought of as analogous to a soaker hose where water flows in at the head of the hose and leaks out through a series of holes. The present method allows the designer to dynamically re-position a portion of

US 6,493,679 B1

11

the leaky feeder antenna and see in real time the effects on wireless system performance at the specified watch points. In the preferred embodiment, distributed antenna systems can be analyzed in terms of the contributions of individual antennas or collections of antennas taken as a whole, providing "composite" results in the latter case.

Referring to FIG. 5, there is shown the general method of the invention. Before one can run an automated predictive model on a desired environment, a 3-D electronic representation of that environment must be created in function block 10. The preferred method for generating a 3-D building or environment database is disclosed in the concurrently filed, copending application Ser. No. 09/318,841, (Docket 256015AA). The resulting definition utilizes a specially formatted vector database format and comprises lines and polygons rather than individual pixels (as in a raster format). The arrangement of lines and polygons in the database corresponds to obstructions/partitions in the environment. For example, a line in a database could represent a wall, a door, tree, a building wall, or some other obstruction/partition in the modeled environment.

From the standpoint of radio wave propagation, each of the obstruction/partition in an environment has several electromagnetic properties. When a radio wave signal intersects a physical surface, several things occur. A certain percentage of the radio wave reflects off of the surface and continues along an altered trajectory. A certain percentage of the radio wave penetrates through or is absorbed by the surface and continues along its course. A certain percentage of the radio wave is scattered upon striking the surface. The electromagnetic properties given to the obstruction/partitions define this interaction. Each obstruction/partitions has parameters that include an attenuation factor, surface roughness, and reflectivity. The attenuation factor determines the amount of power a radio signal loses upon striking a given obstruction. The reflectivity determines the amount of the radio signal that is reflected from the obstruction. The surface roughness provides information used to determine how much of the radio signal is scattered and/or dissipated upon striking an obstruction of the given type.

Once this 3-D database of obstruction data has been built, the design engineer performs computer aided design and experimentation of a wireless network to be deployed in the modeled environment in function block 11, to be described later. Cost and wireless system performance target parameters, transmitters, channel lists, placement options and antenna systems are all taken into account by the present invention.

In order to fine tune the experimental predictions, RF measurements may be optionally taken in function block 12. A preferred method for collecting RF measurements is disclosed in copending application Ser. No. 09/221,985, supra. If necessary, database parameters that define the partition/obstruction characteristics may be modified using RF measurements as a guide to more accurately represent the modeled 3-D environment in function block 13.

The results of the predictive models may be displayed in 3-D overlaid with the RF measurement data, if any, at any time in function block 14. The design engineer analyzes the differences in the predicted and actual measurements in function block 15, and then modifies the RF predictive models, if needed, in function block 16. If necessary, the 3-D environment database may be modified based on the actual measurements to more accurately represent the wireless system coverage areas in function block 10, and so on iteratively until done. The designer can optionally continue with any other step in this process, as desired.

12

The method of invention may be used in a variety of ways depending on the goals of the design engineer. FIG. 6 shows a variant on the above method used to generate estimates based on RF measurements. A 3-D database of the environment must still be generated in function block 10. Field measurements are collected in function block 12. The RF measurement data are then incorporated into the drawing of the environment in function block 61. The design engineer may then generate estimates of power level and location of potential transmitters in function block 62.

FIG. 7 shows a variant of the method used to achieve optimal prediction accuracy using RF measured data. Once again, a 3-D database of the environment is generated in function block 10. However, before collecting field measurements, the design engineer creates a channel plan with "virtual" macrocell locations and power levels in function block 71. The field measurements are then collected in function block 12 and the "virtual" locations of interfering transmitters can be determined in function block 72. The best propagation parameters are then matched with measured data from the interferers in function block 73.

A more detailed description of the method for prediction used in the present invention is now described. Referring to FIG. 8, the 3-D environment definition is input in function block 801. The first step required before predicting the performance of the wireless communication system is to model the wireless system with the 3-D environment. Antennas and types of related components and locations are selected in function block 802. The desired antennas are chosen from a parts list of wireless hardware devices that may include a variety of commercially available devices. Each antenna is placed at a desired location within the environment, for instance, in a specific room on a floor of a building or on a flag pole in front of a building. A number of other components may be created and placed either within or connected to each antenna system. These components include, but are not limited to: cables, leaky feeder antennas, splitters, connectors, amplifiers, or any other user defined component.

FIGS. 9A and 9B show a method for adding antenna systems to a desired environment and generally for running trade-off analyses. First, the designer positions and defines outdoor wireless communication systems, if necessary in function block 901. Next, the designer positions and defines indoor base stations in function block 902. The methods of function blocks 901 and 902 differ in that the components of indoor wireless system will typically be different than an outdoor wireless system. In both cases, the designer is guided through a series of pull down menus and point-and-click options to define the location, type of hardware components and associated performance characteristics of the antenna systems. This data is stored in a database, that also contains cost and manufacturing specific information to produce a complete Bill of Materials list automatically, to be viewed at any time.

In order to fully describe an antenna system in a newly created (or to be modified) wireless system, the designer specifies the air interface/technology and frequencies associated with the wireless system in function block 903. The designer then lays out the full antenna system for the wireless network in function block 904. Components such as base stations, cabling, connectors, amplifiers and other items of the antenna system are then selected from a parts list library containing information on commercially available hardware components in function block 905. Next, the air interface and technology specific parameters are assigned and channel frequencies are customized for the wireless

US 6,493,679 B1

13

system in function block 906. The channel frequencies are selected from pre-assigned channel clusters and assigned to the wireless system in function block 907. An antenna system is then configured in function block 908, selecting antennas from the parts list library as described above. The antennas are placed on the floor plan in function block 909 using a point and click of a mouse or other positioning device to visually place each component in the 3-D database.

At this or any time after a component has been placed on a floor, the designer may view a bill of materials in function block 910. If necessary, the parts list may be modified to add or delete components or modify a component's cost or performance characteristics in function block 911. Components may be replaced or swapped for similar components for a quick trade-off analysis of both wireless system performance and overall cost in function block 912. Components may be added, deleted or modified to more fully define the wireless communications system in function block 913. The designer may redisplay the view of the environment including the wireless communication system, RF measurement data, and/or wireless system predicted performance results in a variety of forms, including 2-D, 3-D wireframe, 3-D wireframe with hidden lines, 3-D shaded, 3-D rendered or 3-D photorealistic rendering, at any time in function block 914.

Typically, a designer will add wireless system components in succession, where each newly placed system component connects to a previously positioned component in the wireless network. One should note that cables and leaky feeder antennas are defined by a series of vertices connected by lines representing lengths of cabling as they are placed on a floor. Cables and leaky feeders may also stretch vertically across building floors, down the sides of buildings, through elevator shafts, etc., simply by adding a vertex in the cable, changing the vertical height, and then continuing to place cable in new locations, in function block 915. The designer does not need to manipulate a 3-D view of the environment and attempt to guide the cables vertically in the 3-D model. The designer may repeat any of the steps in this process, in any order, in the present invention.

Referring again to FIG. 8, once the 3-D environment has been defined and antennas, cables and other objects have been selected and located, the wireless system performance prediction models may be run in function block 803. A variety of different such models are available and may be used in succession, or alone to generate a sufficient number of "what-if" scenarios for predicting and optimizing of antenna placements and component selections.

Referring to FIG. 10, a method for predictive modeling according to the invention is shown. First, the designer selects the desired wireless system performance prediction model in function block 1001. Preferred models are:

Wall/floor Attenuation Factor, Multiple Path Loss Exponent Model,

Wall/floor Attenuation Factor, Single Path Loss Exponent Model,

True Point-to-Point Multiple Path Loss Exponent Model, True Point-to-Point Single Path Loss Exponent Model, Distance Dependent Multiple Breakpoint Model,

Distance Dependent Multiple Path Loss Exponent Model, Distance Dependent Single Path Loss Exponent Model, or other models as desired by the design engineer.

The physical and electrical properties of obstructions in the 3-D environment are set in function block 1002. Although not all parameters are used for every possible predictive model, one skilled in the art would understand

14

which parameters are necessary for a selected model. Parameters that may be entered include:

Prediction configuration—RSSI, C/I, and/or C/N (carrier to noise ratio);

Mobile Receiver (RX) Parameters—power, antenna gain, body loss, portable RX noise figure, portable RX height above floor;

Propagation parameters—

Partition Attenuation Factors

Floor Attenuation Factors

Path Loss Exponents

Multiple Breakpoints

Reflectivity

Surface Roughness

Antenna Polarization

other parameters as necessary for a given model. The designer may save sets of physical, electrical and aesthetic parameters for later re-use. If such a parameter set has been previously saved, the designer may load that set in function block 1003, thereby overwriting any parameters already in selected.

A designer then may select a number of watch points in function block 1004 to monitor for wireless system performance. Referring now to FIG. 11, there is shown a simplified layout of a floor plan with a base station 1100. The designer may use a mouse or other positioning device to point and click to any number of locations in the floor plan to select critical areas, or watch points, for monitoring. Here, for instance, four watch points 1101, 1102, 1103 and 1104 have been selected.

FIG. 12 shows a display, that lists by location, watch points selected for the current prediction. The designer may then select predictions for RSSI, signal to interference ratio (SIR) or signal to noise ratio (SNR). In addition, the designer can see changes in predicted values for each watch point in real time as the mouse is moved, or can choose to select new antenna positions specifically by clicking on a new location. As the designer repositions the mouse cursor, the antenna(s) selected prior to initiating the prediction are effectually repositioned and/or relocated according to position of the cursor. Once all watch points are selected, the prediction model is run. An alternative embodiment is that watch points could be entered and modified on the fly, as the prediction model is being run, rather than defined only prior to running the model. Another alternative embodiment is that RF values at the watch points are updated continuously as the mouse is repositioned, without a click being necessary.

FIG. 13 shows the floor plan of FIG. 11 with the initial RSSI values for each watch point 1101, 1102, 1103 and 1104 also shown. The designer may move the antenna 1100 to a new location and monitor the same watch points for coverage. FIG. 14 shows the floor plan of FIGS. 11 and 13 with the antenna 1100 moved to a new location 1400. The RSSI values at each watch point 1101, 1102, 1103, and 1104 are automatically updated with values associated with the new location of the antenna. Alternatively, the designer may choose to modify the components within the antenna system 1100 for performance or cost reasons. FIG. 15 shows the floor plan of FIGS. 11 and 13 with a base station 1100a at the same location, but with a higher performance antenna component. The RSSI values at each watch point 1101, 1102, 1103, and 1104 are again automatically updated with values associated with the new wireless system performance parameters.

Referring again to FIG. 10, for RF coverage models, the coverage areas and values are displayed in function block 1005. If so desired, the designer modifies the electrical

US 6,493,679 B1

15

parameters of the obstructions, or modified components of antenna systems, or modifies antenna system locations or orientation, etc. in function block 1006 before running another prediction model in function block 1001.

Referring again to FIG. 8, after running a number of models, the design engineer may determine that RF coverage is optimal in decision block 804. If so, then depending on the results either a change in the location of antenna(s) and components will be desired or possibly just a substitution of components without a location change. For instance, even though the coverage may be more than adequate, the total cost of the wireless system could be prohibitive. A method for optimizing the costs using a dynamic, real time, bill of materials management system is disclosed below. Regardless, if the wireless network as currently modeled is not deemed optimal, then the method would continue again in function block 802 to re-select the components.

Once the design is as desired, then the 3-D database holds all of information necessary to procure the necessary components in the Bill of Materials. The locations of each component are clearly displayed, and a visual 3-D representation can be viewed as a guide.

Once the wireless system design is as desired, the database holds all of information necessary to procure the necessary components in the Bill of Materials. The locations of each component are clearly shown, overlaid with the physical environment, and a visual 3-D representation can be viewed as a guide.

Generating and Managing a Bill of Materials

As described above, in more detail, the invention uses 3-D computer aided design (CAD) renditions of a building, collection of buildings, or any other such environment that contains information suitable for the prediction of a wireless system performance. Estimated partition electrical properties can be extracted from radio frequency measurements already published, and/or specified by the designer at any time. Once the appropriate electrical properties are specified, an unlimited number of RF sources can be placed in the 3-D database, and received signal strengths intensity (RSSI) or carrier-to-interference (C/I) ratios can be plotted directly onto the CAD drawing.

The 3-D environment database could be built by a number of methods, the preferred method being disclosed in the concurrently filed, copending application Ser. No. 09/318,841, (Docket 256015AA). Traffic capacity analysis, frequency planning, and Co-channel or adjacent channel interference analysis can be performed concurrently with the prediction of RSSI, C/I and other wireless system performance measures. The antenna system and bill of materials could be built by a number of methods. The preferred method for building the antenna system is described above.

As the designer builds a model of a wireless communications system in a specified environment, as described above, a full bill of materials is maintained for every drawing in the environment. That is, each drawing may contain its own unique set and arrangement of antennas, feed systems and related components representing a variation in the design of a wireless communication system. These components are drawn from a global parts list library. A number of methods could be used to generate the global parts list library, and it would be apparent to one skilled in the art that varying formats could be used.

In the present invention, the design engineer selects a specific wireless system hardware component from the parts list library using pull-down menus and displayed dialog

16

windows. The selection criteria for a particular component is wireless system design dependent, but generally involves the desirability of a component based upon its electrical characteristics and potential effect on wireless system performance, material cost, and/or installation cost. The present invention enables the designer to narrow the focus of component selection to only those devices contained within the parts list library that have the desired characteristics. For example, the design engineer may choose to design a wireless system using components from a specific manufacturer or set of manufacturers that have a desirable material cost and/or electrical characteristics. In doing so, only those devices that meet the requested criteria are displayed for selection from dialog windows in the present invention.

Once a desired component is selected by pointing and clicking with a mouse or other input device, the design engineer may position the component within the three dimension environmental database. This process involves the design engineer using the mouse or other input device to visually identify the desired location for the component by clicking (or otherwise identifying) positions within the 3-D environmental database. For example, an antenna component could be placed within a specific room of a building, atop a flag pole on the side of a building, in the center of a park, or any other location deemed reasonable by the designer. In similar fashion, hardware components that span distances (e.g. coaxial cable, fiber optic cable, leaky feeder antenna, or any component having substantial length) are selected and positioned within the 3-D environment by clicking with the mouse or other input device to identify the vertices (or end points) of the component where each pair of vertices are connected by a time segment representing a portion of cable. Thus, while certain components, such as point antennas or splitters, for example, require only a single point in the 3-D environment to identify placement in the wireless communication system, other components such as distribution cables or distribution antennas require the identification of multiple points joined by line segments to identify placement. In the present invention, unique graphic symbols are utilized to represent each wireless system component and overlaid onto the three-dimensional environmental database enabling the designer to visualize the wireless communication system as it would exist in the physical world. As an example of the graphical display and shown only in two dimensions for convenience, FIG. 4 displays a base station 107 connected via two coaxial cables 402 to two indoor point antennas 403a and 403b.

The present embodiment of the invention provides and links information relating to wireless system component dependence. Such dependencies may include but are not limited to impedance matching of adjoining components, maximum run length, and/or proper termination. Certain components in the parts list library may require pre-existing components to have been positioned within the 3-D environmental database before they themselves may be selected and added to the wireless system. For example, a splitter or other device designed to interconnect two or more independent components may require that an existing component be present in the three dimensional database for the splitter to be connected with. In the previous embodiment of the invention, if the designer chooses to place a hardware component within the 3-D environmental database, and the desired component is dependent upon some other device currently placed in the 3-D database, the designer is prompted through a selection window to identify the dependent component and the selected component is positioned accordingly. In the previous example of the splitter

US 6,493,679 B1

17

component, if the designer chooses to connect the splitter onto the end of an existing cable component by identifying the cable component with the mouse or other input device, the position of the splitter within the three-dimensional database is automatically assigned to be the end of the identified cable. Wireless system components that do not have such dependencies (e.g., base station transceivers) may be freely positioned anywhere within the 3-D environmental database that is deemed suitable by the designer. As this description is specific to one particular implementation, one skilled in the art could see how different implementations could be developed and practiced within the scope of this document.

Using the preferred embodiment of the invention, a designer can model and represent, visually as well as mathematically, complex wireless communication systems involving any number of individual hardware components selected from the parts list library, interconnected with and linked to one another to form complete antenna systems. As each component has associated characteristics regarding electrical properties (e.g. gain, noise figure, attenuation) and cost, the addition, removal, or change of any component directly impacts both the performance of the wireless system and the overall system cost. With the preferred embodiment of the invention, this information is updated in real-time as the designer makes changes to the wireless system. If a wireless communication system includes a specific hardware component, the present invention retrieves the associated electromechanical characteristics and other pertinent information from the parts list library entry that has been specified for the component. This information is stored in a database and is then used to quantify the effect that the component has on various aspects of wireless system design parameters or performance. For example, if the parts list library information for a specific cable indicates that the attenuation loss of the cable is 3.5 dB per 100 meters, and the designer has added a 200 meter segment of the cable to the wireless communication system, the present invention combines the information regarding the placement and length of the cable in the 3-D environmental database with the attenuation loss information from the parts list library to determine a total attenuation loss of 7 dB for the cable. Furthermore, the noise figure and other related qualities of the cable is also computed based upon well known communication theory. If the designer then adds an amplifier to the wireless system and connects it onto the end of the cable as described above, the invention retrieves information regarding the amplifier from the parts list library to determine overall gain of the wireless distribution system. If, for instance, the selected amplifier has an associated gain of 10 dB and some specified noise figure, the present invention combines the characteristics of the interconnected cable and amplifier to determine a total gain of 3 dB for the combined components, and a new system noise figure. If the designer edits or alters component information in the parts list library, this is automatically reflected in the wireless system performance prediction. For example, if the amplifier in the example above has the gain associated with it edited in the parts list library and changed from 10 dB to 15 dB, the combined system characteristics, which may include but are not limited to system gain and system noise figure, of the cable and amplifier from the example are automatically recalculated, resulting in an overall gain of 8 dB instead of 3 dB. Similarly if the cable is repositioned such that its overall length is altered or replaced with a different component from the parts list library, the effect of doing so is automatically recalculated and reflected in all future opera-

18

tions. Although the given example is in terms of simple gains and losses of the individual wireless components, one skilled in the art could apply this same method to any other electrical, electromechanical, financial, aesthetic or other quality associated with components in the parts list library and the overall system in a similar fashion.

A preferred Parts List Library is designed to be generic and applicable to any type of wireless communication system component or wireless communication system design methodology. There are eight basic categories of components in the preferred parts list library utilized in the preferred embodiment, although more categories could be added, as desired:

1. Amplifiers/Attenuators—generally speaking, devices that either boost or decrease the strength of radio wave signals;
2. Connectors/Splitters—generally speaking, devices that connect one or more components to one or more additional components;
3. Cables—various types of cabling (e.g., fiber optic cable, coaxial cable, twisted pair cable, etc);
4. Manufacturer-Specified Point Antennas—any antenna that is manufactured and whose manufacturer has supplied information with regard to the radiation pattern of the antenna. The radiation pattern of an antenna describes the manner in which radio signals are radiated by the antenna. Antenna manufacturers supply radiation pattern information regarding their antennas so that wireless system designers can maximize the effectiveness of antenna deployments;
5. Generic Point Antennas—any generic or idealistic antenna (that is, an antenna that may not be physically realizable or has a generic radiation pattern);
6. Leaky Feeder Cabling/Antennas—a type of antenna that takes the form of a specialized coaxial cable;
7. Base Station/Repeater—the controlling portion of the wireless communication system. The base station manages all communication taking place in the wireless network; and
8. Other—Any component that does not belong in one of the above categories.

Each component has a variety of associated values. These include, but are not limited to:

- Manufacturer Name;
- Manufacturer Part Number;
- User-supplied Description;
- Frequency range at which part has been tested;
- Attenuation/Amplification;
- Number of Connections;
- Physical Cost (material cost of component);
- Installation Cost; and
- Antenna Radiation Pattern.

Base stations and repeater components have a number of additional parameters associated with them, including, but not limited to:

- Technology/Air Interface—identifies the wireless technology employed by the base station (e.g., AMPS ("analog cellular"), IS-136 ("digital cellular"), IEEE 802.11 ("wireless LAN"), etc.);
- Frequency/Channel Assignments—identifies the radio frequencies/channels this base station can utilize; and
- Transmit Power—the amount of power the base station is broadcasting.

US 6,493,679 B1

19

An excerpt from the preferred embodiment of a parts list, with identifying line numbers not actually in the database, is shown below.

- 1: KEY|ITEM|TYPE|MANUFACTURER|PART #|FRE-
QUENCY (MHZ)|dB LOSS/GAIN (per 100 m if
CABLE)|CONNECTIONS|COST (US\$)
|DESCRIPTOR FILE
 - 2: 0|GENERIC FEED LINE|CABLE|GENERIC|N/
A|900|2|2|0|N/A
 - 3: 1|GENERIC
CONNECTOR|CONNECTOR|GENERIC|N/
A|900|1|2|0|N/A
 - 4: 2|GENERIC SPLITTER|CONNECTOR|GENERIC|N/
A|900|2|3|0|N/A
 - 5: 3|GENERIC 10 dB
AMPLIFIER|AMPLIFIER|GENERIC|N/
A|900|10|2|0|N/A
 - 6: 4|GENERIC LEAKY FEEDER|ANTENNA
LEAKY|GENERIC|N/A|1900|4|2|0|N/A
- Line 1 is a header line indicating the titles of fields delimited by a pipe, or "|", character. The first field is the "KEY" field; the second field is the "ITEM" field; the third field is the "TYPE" field; and so on. The next to last field is the cost in U.S. dollars. Lines 2 through 6 show five records of data in the parts list for the following components:
- Generic feed line,
 - Generic connector,
 - Generic splitter,
 - Generic 10 dB amplifier, and
 - Generic leaky feeder.

The parts list can be easily modified by a design engineer as new components are placed on the market, removed from the market or repriced. The ability to maintain a unique equipment list for each drawing enables the designer to carry out rapid design analyses to compare and contrast the performance and cost of different vendor components. The impact of utilizing a specific component in terms of both cost and wireless communication system performance can be seen immediately using the present invention. Information that can be tracked with the bill of materials includes the manufacturer and part number, physical and installation cost, RF loss characteristics, connections, and the frequencies for which the component is valid. In addition, a rich set of customization features is utilized to enable the designer to tailor the parts list library to suit the needs of the target application. Moreover, as components with associated length data, such as cables or leaky feeder antennas, are created, stretched, moved or modified, their associated costs and impact on wireless system performance are automatically updated in the bill of materials to account for the change in length. Furthermore, the parts list is stored as an integral part of the drawing database, allowing the user to recall and archive a system design and all of its particulars. In addition, the wireless communication system performance may be recalculated immediately, using either a standard link budget equation, noise figure equation, or some other metric such as bit error rate or network throughput. This recalculation uses the specific electrical specifications of each component in the system, which is also stored in the bill of materials.

Referring again to the drawings, and more particularly to FIG. 16, there is shown an example of a bill of materials summary for a drawing. A description of the base station "MACROCELL" 1610 is shown to identify the antenna

20

system for which the summary is shown. The first component 1611 is a PCN Panel 1710-1990 92 Deg 9.00 dB Gain point antenna manufactured by Allen Telecom. One should note that the component cost 1612, sub-total cost 1613 and total system cost 1614 is \$0.00. This shows that the designer has not yet updated the parts list library with current costs. When the list has been updated, the summary will automatically show component costs as well as sub-totals and totals for all base stations and components in the drawing.

FIG. 17 show a bill of materials where costs have been entered into the parts list database. Another component 1720 has been added to the "MACROCELL" base station, also. The costs of each component 161 2a and 1721 are now shown. Sub-total 1613a and Total costs 1614a are also shown.

Referring now to FIG. 18, the general method of the invention is shown. As previously described, first the designer must create a database defining the desired environment in function block 180. A preferred method being disclosed in the concurrently filed, copending application Ser. No. 09/318,841, (Docket 256015AA). A database of components is then developed in function block 181. In the case of wireless communication networks, a preferred method is described above. The creation of these components will automatically generate a parts list categorized by base station and antenna system. A bill of materials may be displayed at any time in function block 182.

In order to optimize the design of the wireless communications system and ensure adequate antenna coverage, the designer runs a series of prediction models and optimization techniques in function block 183. A preferred method for running predictions is described above. This method allows the designer to see, in real-time, changes in coverage, generally, and for specifically chosen watch points, as antennas are repositioned or reoriented. The designer may choose to add, delete or substitute components in function block 184 and then re-run the models again in function block 183. Each time the designer makes a modification in the system to improve performance, the bill of materials is automatically updated. The designer may run the prediction models in function block 183, and determine if the wireless system, as designed, is adequate in terms of performance and cost. If not, the designer can choose to modify components using cost or component performance considerations. Performance parameters may be entered to enable the designer to choose substitute components from a list that contains only those components that would not degrade the performance of the overall system. Note that in the preferred embodiment, the prediction or system performance models are recomputed upon user demand, but that it would be apparent to one skilled in the art to also have models recomputed instantly ("on-the-fly") as new components are added or subtracted from the bill of materials.

The integration of the bill of materials and component performance specifications is key to providing a quick and efficient method to design high performance wireless communication networks that are within budget.

While the invention has been described in terms of its preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

Having thus described our invention, what we claim as new and desire to secure by Letters Patent is as follows:

I. An automated method for generating a bill of materials for a wireless communication environment, said method comprising the steps:

- (a) inputting data defining an environment in which a system of components is to be built;

US 6,493,679 B1

21

- (b) inputting data defining a plurality of system components, each said system component associated with performance and cost data;
- (c) running prediction models using said data defining an environment and said data defining a plurality of system components to predict performance characteristics of said system of components; and
- (d) generating a bill of materials containing cost information based on cost data associated with said system components.

2. A method as recited in claim 1, wherein said step of inputting data defining an environment is performed using an input device to point and click on a screen display of a three-dimensional representation of said environment in which a system of components is to be built.

3. A method as recited in claim 1, wherein said step of inputting data defining a plurality of system components is performed using an input device to point and click on a screen display of a three-dimensional representation of said environment in which a system of components is to be built.

4. A method as recited in claim 1, wherein said cost information is generated for components, sub-systems and an overall system.

5. A method as recited in claim 1, wherein said generating step generates a bill of materials wherein each said component in said bill of materials is associated with performance characteristics.

6. A method as recited in claim 1, wherein the step of inputting data defining a plurality of system components, further comprises the step of selecting a plurality of system components from a pre-defined list of system components, wherein each said system component in said pre-defined list is associated with performance and cost data.

7. A method as recited in claim 6, wherein the step of selecting a plurality of system components from a pre-defined list of system components further comprises the steps:

- selecting at least one category of system components;
- selecting a range of acceptable performance criteria for each said at least one category of system components; and
- selecting at least one system component in said at least one category of system components from a pre-defined list of system components, wherein system components not meeting said range of acceptable performance criteria are automatically removed from said pre-defined list of system components, thereby allowing selection only of system components which meet said range of acceptable performance criteria.

8. A method as recited in claim 6, wherein the step of selecting a plurality of system components further comprises the step of storing said data defining a plurality of system components in a database, wherein each said system component is associated with a location in a three-dimensional representation of said environment.

9. A method as recited in claim 1, further comprising the step of modifying said data defining a plurality system components, and updating said bill of materials generated in step (d).

10. A method as recited in claim 9, wherein said step of updating said bill of materials is performed automatically when modifying said data defining a plurality system components.

11. A method as recited in claim 9, wherein the step of modifying a set of data includes repositioning at least one said system component in said environment, reorienting at

22

least one said system component in said environment, or substituting at least one said system component for a similar component with different cost or performance characteristics.

12. A method as recited in claim 11, wherein repositioning at least one said system component is performed using an input device to point and click on a screen display of a three-dimensional representation of said environment in which a system of components is to be built in order to quickly assess the changed performance characteristics of said system of components.

13. An apparatus for generating a bill of materials for a wireless environment, comprising:

means for representing and inputting data defining an environment in which a system of components is to be built;

means for representing and inputting data defining system components, each said system component associated with performance and cost data;

means for predicting performance characteristics of said system of components using said data defining an environment and said data defining system components; and

means for automatically generating a bill of materials containing cost information based on cost data associated with system components.

14. An apparatus as recited in claim 13, further comprising a means for modifying said data defining system components, and updating said bill of materials generated by said generating means.

15. An apparatus as recited in claim 14, wherein said means for representing and inputting data defining a plurality of system components, further comprises a means for selecting a plurality of system components from a pre-defined list of system components, wherein each said system component in said pre-defined list is associated with performance and cost data.

16. An apparatus as recited in claim 15, wherein said means for representing and inputting data defining a plurality of system components, further comprises a means for storing said data defining a plurality of system components in a database, wherein each said system component is associated with a location in a three-dimensional representation of said environment.

17. An apparatus as recited in claim 15, wherein said means for selecting a plurality of system components from a pre-defined list of system components selects at least one system component in at least one category of system components from a pre-defined list of system components, wherein system components not meeting a desired range of acceptable performance criteria are automatically removed from said pre-defined list of system components, thereby allowing selection only of system components which meet said range of acceptable performance criteria.

18. An apparatus as recited in claim 14, wherein said means for modifying a set of data repositions at least one said system component, reorients at least one said system component, or substitutes at least one said system component for a similar component with different cost or performance characteristics.

19. An apparatus as recited in claim 18, wherein said means for modifying a set of data repositions at least one said system component by utilizing an input device, to point and click on a screen display of a three-dimensional representation of said environment in which a system of components is to be built.

* * * * *

EXHIBIT 15




DATA SHEET

Aruba AP-70 Access Point

ARUBA AP-70 ACCESS POINT

The Aruba AP-70 is a dual-radio (dual-band concurrent 802.11a plus b/g) high-performance indoor wireless access point capable of supporting a wide range of functions including WLAN access, air monitoring/wireless intrusion detection and prevention, and secure enterprise mesh across the 2.4-2.5 GHz and 5 GHz RF spectrums. The AP-70 access point supports diverse deployment options, delivering secure user-centric network services and applications in enterprise and campus environments, branch offices, retail spaces, and to remote locations over public or private networks as an advanced featured Remote AP. Centrally managed from an Aruba mobility controller, the AP-70 empowers the network administrator with unparalleled control over services, security and deployment models. The AP-70 supports unparalleled interface flexibility in an access point, supporting dual 10/100 Ethernet Interfaces, redundant 802.3af PoE sourcing, and a USB 2.0 interface for service extension. The AP-70 features onboard dual, integral, high-performance omni-directional multi-band antennas and also supports external antennas through quad, detachable antenna interfaces.



APPLICATION

- High-performance enterprise and campuses, branch offices and retail spaces where flexibility in interfaces is required. Remote access and secure-jack applications. Indoor applications.

OPERATING MODE

- Multi-service 802.11a+b/g WLAN, 802.11a+b/g air monitor, hybrid combination of WLAN/AM and Remote AP, all with secure jack.

RADIOS

- Dual radio - software configurable to 802.11a and 802.11b/g

RF MANAGEMENT

- Automatic transmit power and channel management control with auto coverage hole correction via ARM

MOBILITY SERVICE DELIVERY

- Virtual AP Services:
 - Supports up to 32 SSIDs per access point
 - Multiple captive portals per SSID
 - Support any combination of encryption / authentication types per SSID
 - Session level QoS
 - VLAN load balancing
 - Guest account creation / management
- Voice Services:
 - Wireless Multi-media QoS (WMM)
 - 802.1p and DSCP to WMM AC tagging
 - Upstream traffic prioritization
 - Call Admission Control (CAC)
 - Traffic classification/session bandwidth reservation (T-SPEC TCLAS)
 - Unscheduled power save delivery (U-APSD)
 - Stateful session awareness (soft voice client QoS):
 - SIP
 - NOE
 - Cisco Skinny
 - Vocera

- Spectralink Voice Prioritization (SVP)
- Support for proxy-ARP and multicast filtering
- Battery Boost
- Priority queuing
- Voice-aware scanning support in ARM

802.11A RADIO SPECIFICATIONS

- Operating Frequency: 5.150 GHz - 5.950 GHz*
- Available Channels: Mobility controller-managed, dependent upon configured regulatory domain
- Modulation: Orthogonal Frequency Division Multiplexing (OFDM)
- Transmit Power: Configurable in increments of 0.5 dBm
- Association Rates (Mbps):
 - 54, 48, 36, 24, 18, 12, 9, 6 with automatic fallback

802.11B RADIO SPECIFICATIONS

- Operating Frequency: 2.4 GHz - 2.5 GHz
- Available Channels: Mobility controller-managed, dependent upon configured regulatory domain
- Modulation: Direct-Sequence Spread-Spectrum (DSSS)
- Transmit Power: Configurable in increments of 0.5 dBm
- Association Rates (Mbps): 11, 5.5, 2, 1 with automatic fallback

802.11G RADIO SPECIFICATIONS

- Operating Frequency: 2.4 GHz - 2.5 GHz
- Available Channels: Mobility controller-managed, dependent upon configured regulatory domain
- Modulation: Orthogonal Frequency Division Multiplexing (OFDM)
- Transmit Power: Configurable in increments of 0.5 dBm
- Association Rates (Mbps): 54, 48, 36, 24, 18, 12, 9, 6 with automatic fallback

AVAILABLE CHANNELS 802.11A/B/G

- Centrally managed by mobility controller, based on configured regulatory domain
- Mobility controller-managed, dependent upon configured regulatory domain

ARUBA AP-70 ACCESS POINT**ANTENNA**

- Integral, dual, omni-directional multi-band dipole (supports spatial diversity)
 - Gain:
 - 2.4 GHz-2.5 GHz / 4.46 dBi
 - 5.150 GHz / 7.21 dBi
 - 5.350 GHz / 6.49 dBi
 - 5.850 GHz / 5.23 dBi
- Quad, RP-SMA interfaces (2 per radio) for external antennas

INTERFACES

- Network:
 - 2 x 10/100Base-T Ethernet (RJ45) , Auto-sensing link speed and MDI/MDX
 - 48 V DC IEEE compliant 802.3af Power-over-Ethernet (PoE) (load sharing over both ports)
 - Serial-over-Ethernet (SoE) (primary port)
- Power:
 - 1 x 5 V DC up to 2.5 A (for external AC adapter power)
- Antenna:
 - 4 x RP-SMA antenna interfaces (2 per radio)

POWER

- 48 V DC IEEE compliant 802.3af Power-over-Ethernet (PoE)
- 5 V DC for external AC supplied power (adapter sold separately)

MOUNTING

- Standard:
 - Desk-top
 - Wall
- Optional mounting kit:
 - Secure wall
 - Ceiling tile rail (15/16")
- Security:
 - Security screw
 - Kensington security lock point

MECHANICAL

- Dimensions/Weight:
- Antenna Stowed:
 - 7.4" x 6.8" x 1.4"
 - 188 mm x 173 mm x 36 mm
 - 1.15 lbs/0.52 Kgs
- Antenna Extended 180°:
 - 7.4" x 11.7" x 1.4"
 - 188 mm x 295 mm x 36 mm
- Dimensions/Weight (Shipping):
 - 10.1" x 10.4" x 4.1"
 - 257 mm x 264 mm x 104 mm
 - 2.0 lbs/0.91 Kgs

ENVIRONMENTAL

- Operating:
 - Temp: 0° to 50° C (32° to 122° F)
 - Humidity: 5 to 95% non-condensing
- Storage:
 - Temp: 0° to 70° C (32° to 158° F)

REGULATORY

- FCC Part 15
- Industry of Canada
- VCCI
- MIC
- PSE mark - adapters/cords
- Anatel
- NOM/COFETEL
- SRRC
- GS Mark
- CE Mark
- R&TTE Directive - 1995/5/EC
- Low Voltage Directive - 72/23/EEC
- EN 300 328
- EN 301 893
- EN 301 489
- UL/IEC/EN 60950-1:2001 CB, cULus
- AS/NZS 4268, 4771
- Medical EN 60601-1, -2
- UL2043 Listed

(For more country-specific regulatory information, and approvals, please see your Aruba representative)

CERTIFICATIONS

- Wi-Fi Certified: 802.11a/b/g

WARRANTY

- 90 days parts/labor

PART NUMBERS

- Access Point:
 - AP-70
- Accessories:
 - AP-70-MNT (mount kit)
 - AP-ANT-1 to 13 (detachable antennas)
- AC Power Adapters:
 - AP-AC-NA (North America)
 - AP-AC-JPN (Japan)
 - AP-AC-UK (UK)
 - AP-AC-EC (Central Europe)
 - AP-AC-IT (Italy)



WWW.ARUBANETWORKS.COM


1322 Crossman Avenue, Sunnyvale, CA 94089 | Tel. +1 408.227.4500 | Fax. +1 408.227.4550

EXHIBIT 16



AP300 Access Port

Next-generation access ports:
more functionality at a lower cost



FEATURES

Dual form-factors

Plenum-rated external antenna model with metal housing is ideal for installation above ceiling tiles; the plastic internal antenna housing allows for installation within the "carpeted space" and provides cost-effective coverage via the integrated 2.4 GHz and 5.2 GHz antennas.

Interoperability

Standards-based wireless, wired and security protocols ensure interoperability with third-party hardware.

802.1x supplicant

Allows authentication to a RADIUS server to enable an 802.1x-protected Ethernet port.

802.11h

Enables worldwide operation through support for standards-based dynamic frequency selection and power control.

More functionality for a fraction of the cost of access points

Access ports are a key component of Motorola's award winning wireless switch system, the wireless LAN architecture that does more, yet costs less. Working in conjunction with Motorola's wireless switches, the AP300 Access Port delivers robust and feature rich IEEE 802.11a/b/g connectivity. It can also be used as a sensor in conjunction with Motorola's Wireless Intrusion Protection System (IPS). Access ports substantially reduce the cost of deploying, implementing and managing a wireless LAN, while significantly increasing features, functionality and security of the wireless LAN infrastructure.

Virtual AP enables true RF Virtual LANs (VLANs) for better device and network performance

With Virtual AP, each access port can support four separate wireless broadcast domains — functionality that would otherwise require the installation of four first-generation access points. These true wireless VLANs enable separation of mobile end-users, ensuring that broadcast traffic reaches only those recipients for whom it is intended. Overall network traffic is reduced, network and device performance is improved, and device battery life is increased — at a fraction of the cost required to deliver the same functionality

in a first generation access point-based network. Each AP300 supports four BSSIDs (Basic Service Set Identifiers) and 16 ESSIDs (Extended Service Set Identifiers) per radio, enabling granular segmentation of the wireless LAN into multiple broadcast domains to meet specific enterprise needs. Typical access points support only one BSSID, utilizing ESSIDs (instead of BSSIDs) to create wireless VLANs.

Dual-radio 802.11a and 802.11g design

Simultaneous service to 802.11a, 802.11b and 802.11g mobile devices provides high-bandwidth wireless connectivity at speeds of up to 54 Mbps in both the 2.4 GHz and 5.2 GHz ISM bands.

Thin AP design

The AP300, as all other Motorola access ports, requires no configuration or manual firmware maintenance. The Motorola wireless switch discovers access ports on the network and automatically downloads all configuration parameters and firmware, greatly reducing installation, maintenance and troubleshooting costs.

For more information, contact Motorola at +1.800.722.6234 or +1.631.738.2400, or visit us on the web at: motorola.com/ap300

SPECIFICATION SHEET

AP300 ACCESS POINT

Next generation access ports. More functionality at a lower cost.

802.11i

Support for IEEE standards-based security protocols for strong Encryption (AES, TKIP), Authentication and Key Management (802.1x-EAP)

Flexible mounting options

Fast and easy installation with wall, ceiling and above-ceiling tile mounting options; internal antenna version snaps on to T-bars of suspended ceilings without the use of any hardware; external antenna version installs above ceiling tiles

802.3af

Simplifies and reduces total cost of installation through support of standards-based Power over Ethernet (PoE)

Load balancing, pre-emptive roaming and rate scaling

Increases reliability and resilience of the wireless network to support mission critical applications

AP300 Specifications

Physical Characteristics	AP300 (internal antenna)	AP300 (external antenna)
Dimensions:	9.5 in. L x 7.0 in. W x 2.0 in. H/24.1 cm L x 17.8 cm W x 5.1 cm H	9.25 in. L x 5.75 in. W x 1.0 in. H/23.5 cm L x 14.6 cm W x 2.54 cm H
Weight:	1.0 lbs./0.45 kg	1.6 lbs./0.73 kg
Part Number.*	WSAP-5110-100-WWR; WSAP-5110-050-WWR	WSAP-5100-100-WWR; WSAP-5100-050-WWR
Available Mounting Configurations:	Ceiling-mount (to suspended ceiling T-bars, below tile); wall mount	Ceiling-mount (above tile); wall-mount
Plenum Rated:	No	Yes, certified to UL 2043
LEDs Indicators:	2 LED indicators with multiple modes indicating 802.11a/802.11g Activity, Power, Adoption and Errors	
Wireless Data Communications		
Data Rates Supported:	802.11a: 6, 9, 12, 18, 24, 36, 48 and 54 Mbps; 802.11b/g: 1, 2, 5.5, 6, 9, 11, 12, 18, 24, 36, 48, 54 Mbps	
Network Standard:	802.11a, 802.11b, 802.11g	
Wireless Medium:	Direct Sequence Spread Spectrum (DSSS) and Orthogonal Frequency Division Multiplexing (OFDM)	
Uplink:	Auto-sensing 10/100Base-T Ethernet	
Radio Characteristics		
Frequency:	802.11b/g: 2.412 GHz to 2.484 GHz; 802.11a: 4.9 GHz to 5.875 GHz	
FCC (US and Canada):	2.412 GHz to 2.462 GHz; 5.150 GHz to 5.250 GHz (UNII -1); 5.250 GHz to 5.350 GHz (UNII -2); 5.725 GHz to 5.825 GHz (UNII -3); 5.825 GHz to 5.850 GHz (ISM)	
EU:	2.412 GHz to 2.472 GHz; 5.150 GHz to 5.250 GHz; 5.150 GHz to 5.350 GHz; 5.470 GHz to 5.725 GHz; (Country Specific)	
Japan:	2.412 GHz to 2.484 GHz; 4.900 GHz to 5.000 GHz; 5.150 GHz to 5.250 GHz	
China:	2.412 GHz to 2.472 GHz	5.725 GHz to 5.850 GHz
Operating Channels:	802.11b/g: ETSI: 13; North America: 11; TELEC (Japan): 13 802.11a: ETSI: Country Specific; North America: 12; UNII I, II, III; (approval for 5.4-5.7 GHz pending); TELEC (Japan): 8	
Nominal Transmitter Power:	802.11b/g: 17.5 dBm +/- 1 dBm @ 1, 2, 5.5, 11 Mbps; 17.0 dBm +/- 1 dBm @ 6 and 9 Mbps; 16.5 dBm +/- 1 dBm @ 12 and 18 Mbps; 14.0 dBm +/- 1 dBm @ 24 and 36 Mbps; 12.5 dBm +/- 1 dBm @ 48 and 54 Mbps 802.11a: 17.5 dBm +/- 1 dBm @ 6 and 9 Mbps; 16.0 dBm +/- 1 dBm @ 12 and 19 Mbps; 14.0 dBm +/- 1 dBm @ 24 and 36 Mbps; 12.0 dBm +/- 1 dBm @ 48 and 54 Mbps	
Receiver Sensitivity:	802.11b: 11 Mbps @ -84dBm; 5.5 Mbps @ -87dBm; 2 Mbps @ -88dBm; 1 Mbps @ -90dBm 802.11g: 54 Mbps @ -68 dBm; 48 Mbps @ -70 dBm; 36 Mbps @ -75 dBm; 24 Mbps @ -79 dBm; 18 Mbps @ -81 dBm; 12 Mbps @ -85 dBm; 9 Mbps @ -87 dBm; 6 Mbps @ -88 dBm 802.11a: 54 Mbps @ -68 dBm; 48 Mbps @ -70 dBm; 36 Mbps @ -75 dBm; 24 Mbps @ -79 dBm; 18 Mbps @ -81 dBm; 12 Mbps @ -85 dBm; 9 Mbps @ -87 dBm; 6 Mbps @ -88 dBm	
User Environment		
Operating Temperature:	32°F to 104° F/0°C to 40° C	-4°F to 122° F/-20°C to 50° C
Storage Temperature:	-40°F to 158° F/-40°C to 70° C	
Operating Humidity:	5%-95% (non-condensing)	
Operating Altitude:	8,000 ft./2438 m	
Storage Altitude:	15,000 ft./4572 m	
Electrostatic Discharge:	+/- 15 kV (Air), +/- 8 kV (Contact)	
Power Specifications		
Operating Voltage:	48 VDC @ 7W (Typical), 36 VDC to 57 VDC (Range)	
Operating Current:	145mA @ 48VDC (typical)	
Integrated Power-over-Ethernet Support:	Standards-based IEEE 802.3af	
Antenna Specifications		
Type:	Integrated 2.4 GHz and 5.2 GHz Dual-Antenna, Elements with diversity	Two RSMA and two RBNC connectors for external antennas (not included)
Band:	2.4 GHz to 2.5 GHz; 4.9 GHz to 5.850 GHz (actual operating frequencies depend on regulatory rules and certification agency)	
VSWR:	2.4 GHz: Less than 2:1; 5.2 GHz: Less than 1.5:1	(antenna-specific)
Gain:	2.4 GHz: 0.0 dBi; 5.2 GHz: 3.0 dBi	(antenna-specific)
Regulatory		
Product Safety Certifications:	UL 60950, cUL, EU EN 60950, TUV and UL 2043 (external antenna)	
Radio Approvals:	FCC (USA), Industry Canada, CE (Europe) and TELEC (Japan)	
*WS-2000-1C-ABG-WWR (WS2000 and 1 AP300 (802.11a/b/g) bundle); WS-2000-2C-ABG-WWR (WS2000 and 2 AP300 (802.11a/b/g) bundle); WS-2000-2C-BG-WWR (WS2000 and 2 AP300 (802.11b/g) bundle)		


*WS-2000-1C-ABG-WWR (WS2000 and 1 AP300 (802.11a/b/g bundle); WS-2000-2C-ABG-WWR (WS2000 and 2 AP300 (802.11a/b/g bundle); WS-2000-2C-BG-WWR (WS2000 and 2 AP300 (802.11b/g bundle)

**MOTOROLA**

motorola.com

Part number SS-AP300. Printed in USA 05/07. MOTOROLA and the Stylized M Logo and Symbol and the Symbol Logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. ©2007 Motorola, Inc. All rights reserved. For system, product or services availability and specific information within your country, please contact your local Motorola office or Business Partner. Specifications are subject to change without notice.

EXHIBIT 17



DATA SHEET Aruba MC-2400 Mobility Controller

ARUBA MC-2400 MOBILITY CONTROLLER

The Aruba MC-2400 Mobility Controller is a fully-featured wireless LAN controller that aggregates up to 48 controlled access points (APs) and delivers centralized control and security for wireless deployments. The MC-2400 provides a truly user-centric network experience, delivering follow-me connectivity, identity-based access and application continuity services for regional headquarters or dense office wireless deployments. It can be easily deployed as an overlay without any disruption to the existing wired network and centrally managed using ArubaOS or the Aruba Mobility Management System. Advanced convergence features such as Call Admission Control (CAC), voice-aware RF management and strict over-the-air QoS allow the MC-2400 to deliver mobile VoIP capabilities.



In addition, the MC-2400 can be deployed as an identity-based security gateway to authenticate wired and wireless users, enforce role-based access control policies and quarantine unsafe endpoints from accessing the corporate network. Guest users can easily and safely be supported with the built-in captive portal server and advanced network services. The MC-2400 can create a secure networking environment without requiring additional VPN/firewall devices using its integrated site-to-site VPN and NAT capabilities, split-tunneling and an ICSA-compliant stateful firewall. Site-to-site VPN support can be integrated with all leading VPN concentrators to provide seamless integration into existing corporate VPNs.

CONTROLLER PERFORMANCE AND CAPACITY

Controlled APs	48
Users	768
MAC addresses	4,096
VLAN IP interfaces	128
Fast Ethernet ports (10/100)	24
Gigabit Ethernet ports (GBIC)	2
Active firewall sessions	64,000
Concurrent IPsec tunnels	768
Firewall throughput	2 Gbps
Encrypted throughput (3DES and AES-CCM)	400 Mbps

WIRELESS LAN SECURITY AND CONTROL FEATURES

- 802.11i security (WFA-certified WPA2 and WPA)
- 802.1X user and machine authentication
- EAP-PEAP, EAP-TLS, EAP-TTLS support
- Centralized AES-CCM, TKIP and WEP encryption
- 802.11i PMK caching for fast roaming applications
- EAP offload for AAA server scalability and survivability
- Stateful 802.1X authentication for standalone APs
- MAC address, SSID and location-based authentication
- Multi-SSID support for operation of multiple WLANs
- SSID-based RADIUS server selection
- Secure AP control and management over IPsec or GRE
- CAPWAP-compatible and upgradeable
- Distributed WLAN mode for remote AP deployments
- Simultaneous centralized and distributed WLAN support

IDENTITY-BASED SECURITY FEATURES

- Wired and wireless user authentication
- Captive portal, 802.1X and MAC address authentication
- Username, IP address, MAC address and encryption key binding for strong network identity creation
- Per-packet identity verification to prevent impersonation
- Endpoint posture assessment, quarantine and remediation
- Microsoft NAP, Cisco NAC, Symantec SSE support
- RADIUS and LDAP-based AAA server support
- Internal user database for AAA server failover protection
- Role-based authorization for eliminating excess privilege
- Robust policy enforcement with stateful packet inspection
- Per-user session accounting for usage auditing
- Web-based guest enrollment with Aruba GuestConnect™
- Configurable acceptable use policies for guest access
- XML-based API for external captive portal integration
- xSec option for wired LAN authentication and encryption (802.1X authentication, 256-bit AES-CBC encryption)

CONVERGENCE FEATURES

- Voice and data on a single SSID for converged devices
- Flow-based QoS using Voice Flow Classification™
- SIP, Spectralink SVP, Cisco SCCP and Vocera ALGs
- Strict priority queuing for over-the-air QoS
- 802.11e support – WMM, U-APSD and T-SPEC
- QoS policing for preventing network abuse via 802.11e
- Diffserv marking and 802.1p support for network QoS
- On-hook and off-hook VoIP client detection
- VoIP call admission control (CAC) using VFC
- Call reservation thresholds for mobile VoIP calls

ARUBA MC-2400 MOBILITY CONTROLLER

- Voice-aware RF management for ensuring voice quality
- Fast roaming support for ensuring mobile voice quality
- SIP early media and ringing tone generation (RFC 3960)
- Per-user and per-role rate limits (bandwidth contracts)

ADAPTIVE RADIO MANAGEMENT™ (ARM) FEATURES

- Automatic channel and power settings for controlled APs
- Simultaneous air monitoring and end user services
- Self-healing coverage based on dynamic RF conditions
- Dense deployment options for capacity optimization
- AP load balancing based on number of users
- AP load balancing based on bandwidth utilization
- Coverage hole and RF interference detection
- 802.11h support for radar detection and avoidance
- Automated location detection for active RFID tags
- Built-in XML-based Location API for RFID applications

WIRELESS INTRUSION PROTECTION FEATURES

- Integration with WLAN infrastructure
- Simultaneous or dedicated air monitoring capabilities
- Rogue AP detection and built-in location visualization
- Automatic rogue, interfering and valid AP classification
- Over-the-air and over-the-wire rogue AP containment
- Adhoc WLAN network detection and containment
- Windows client bridging and wireless bridge detection
- Denial of service attack protection for APs and stations
- Misconfigured standalone AP detection and containment
- 3rd party AP performance monitoring and troubleshooting
- Flexible attack signature creation for new WLAN attacks
- EAP handshake and sequence number analysis
- Valid AP impersonation detection
- Frame floods, Fake AP and Airjack attack detection
- ASLEAP, death broadcast, null probe response detection
- Netstumbler-based network probe detection

STATEFUL FIREWALL FEATURES

- Stateful packet inspection tied to user identity or ports
- Location and time-of-day aware policy definition
- 802.11 station awareness for WLAN firewalling
- Over-the-air policy enforcement and station blacklisting
- Session mirroring and per-packet logs for forensic analysis
- Detailed firewall traffic logs for usage auditing
- ICSA corporate firewall 4.1 compliance
- Application Layer Gateway (ALG) support for SIP, SCCP, RTSP, Vocera, FTP, TFTP, PPTP
- Source and destination Network Address Translation (NAT)
- Dedicated flow processing hardware for high performance
- TCP, ICMP denial of service attack detection and protection
- Policy-based forwarding into GRE tunnels for guest traffic
- External service interface for 3rd-party security integration for inline anti-virus, anti-spam and content filtering apps
- Health checking and load balancing for external services

VPN SERVER FEATURES

- Site-to-site VPN support for branch office deployments
- Site-to-site interoperability with 3rd-party VPN servers
- VPN server emulation for easy integration into WLAN

- L2TP/IPsec VPN termination for Windows VPN clients
- Mobile edge client shim for roaming with RSA Tokens
- XAUTH/IPsec VPN termination for 3rd-party clients
- PPTP VPN termination for legacy VPN integration
- RADIUS and LDAP server support for VPN authentication
- PAP, CHAP, MS-CHAP and MS-CHAPv2 authentication
- Hardware encryption for DES, 3DES, AES, MPPE
- Secure point-to-point xSec tunnels for L2 VPNs

NETWORKING FEATURES AND ADVANCED SERVICES

- L2 and L3 switching over-the-air and over-the-wire
- VLAN pooling for easy, scalable network designs
- VLAN mobility for seamless L2 roaming
- Proxy mobile IP and proxy DHCP for L3 roaming
- Built-in DHCP server and DHCP relay
- VRRP-based N+1 controller redundancy (L2)
- AP provisioning-based N+1 controller redundancy (L3)
- Wired access concentrator mode for centralized security
- Etherchannel support for link redundancy
- 802.1d Spanning Tree Protocol (STP)
- 802.1Q VLAN tags

CONTROLLER-BASED MANAGEMENT FEATURES

- RF Planning and AP Deployment Toolkit
- Centralized AP provisioning and image management
- Live coverage visualization with RF heat maps
- Detailed statistics visualization for monitoring
- Remote packet capture for RF troubleshooting
- Interoperable with Ethereal and Airopoke analyzers
- Multi-controller configuration management
- Location visualization and device tracking
- System-wide event collection and reporting

CONTROLLER ADMINISTRATION FEATURES

- Web-based user interface access over HTTP and HTTPS
- Quickstart screens for easy controller configuration
- CLI access using SSH, Telnet and console port
- Role-based access control for restricted admin access
- Authenticated access via RADIUS, LDAP or Internal DB
- SNMPv3 and SNMPv2 support for controller monitoring
- Standard MIBs and private enterprise MIBs
- Detailed message logs with syslog event notification

CONTROLLER POWER SPECIFICATIONS

Power consumption	300 Watts
AC input voltage	90~132/180~264 VAC
AC input frequency	47-63 Hz

OPERATING SPECIFICATIONS AND DIMENSIONS

Operating temperature range	0° to 40° C
Storage temperature range	10° to 70° C
Humidity, non-condensing	5 to 95%
Height	1.75" (44.5 mm)
Width	17.4" (444 mm)
Depth	16.1" (409 mm)
Weight	12.5 lbs. (unboxed)

ARUBA MC-2400 MOBILITY CONTROLLER**WARRANTY**

Hardware	1 year parts/labor*
Software	90 days*

REGULATORY AND SAFETY COMPLIANCE

- FCC part 15 Class A CE
- Industry Canada Class A
- VCCI Class A (Japan)
- EN 55022 Class A (CISPR 22 Class A), EN 61000-3,
- EN 61000-4-2, EN 61000-4-3, EN 61000-4-4,
- EN 61000-4-5, EN 61000-4-6, EN 61000-4-8,
- EN 61000-4-11, EN 55024, AS/NZS 3548
- UL 60950, EN60950
- CAN/CSA 22.2 #60950
- CE mark, cTUVus, GS, CB, C-tick, Anatel, NOM, MIC, IQC

ORDERING INFORMATION

PART NUMBER	DESCRIPTION
2400-48-AOS-STD	A2400-48 Mobility Controller - SPOE, GBIC GigE - 48 AP Limit
LC-GBIC-T	Aruba GBIC Interface Adapter - T
LC-GBIC-SX	Aruba GBIC Interface Adapter - SX
LC-GBIC-LX	Aruba GBIC Interface Adapter - LX

Please contact your Aruba Networks sales representative for more information on configuring and ordering this product

* Extended with support contract



WWW.ARUBANETWORKS.COM

1322 Crossman Avenue, Sunnyvale, CA 94089 | Tel. +1 408.227.4500 | Fax. +1 408.227.4550


EXHIBIT 18

SPECIFICATION SHEET



WS5100

High performance overlay wireless switch



FEATURES

Centralized architecture

A single point of entry that can be centrally managed, easily secured, and lowers the overall cost of deployment and management.

Adaptive AP — Extending the Enterprise

Enables centralized management of mesh access points at remote sites as well as site survivability of those remote locations.

L2 and L3 roaming

Supports inter-subnet roaming without additional client support, seamless roaming of mobile clients within a simplified wireless network design.

Security

Comprehensive, layered security capabilities including WPA2-CCMP (with 802.11i fast roaming options), integrated RADIUS Server, IPSec VPN Gateway, Secure Guest Access Provisioning and advanced wireless intrusion detection, exceptional level of data and network protection without sacrificing fast roaming.

Moving at the speed of business

The WS5100 Wireless Switch from Motorola provides enhanced support for enterprise mobility and multimedia applications, as well as increased security and manageability. Based on Motorola's Wi-NG (Wireless Next Generation) architecture, the WS5100 enables seamless campus-wide roaming, more robust failover capabilities, enhanced security, improved mobile client battery life, and increased voice capacity. Robust security features include integrated intrusion detection, an IPSec VPN gateway, and secure guest access provisioning. Automatic configuration and firmware updates, built-in process monitors, troubleshooting tools and a simple user interface make network deployment and management easy.

Robust enterprise mobility

Business needs should dictate network coverage, not the other way around. That's why the Motorola WS5100 allows you to deploy "thin" access ports in Layer 3 network designs, and enables campus-wide roaming of mobile clients across Layer 3 boundaries — without requiring additional client software or hardware. Used in concert with Motorola handheld devices, the WS5100 further enhances the fast roaming capabilities and extends client battery life. WMM (Wi-Fi Multimedia) with "power save" extensions also provides additional voice capacity. Supporting mobile workers has never been so easy. With the adaptive AP support, Enterprise Mobility is taken a step further, providing capabilities to extend the enterprise to remote/branch locations with the advantages of centralized manageability.

End-to-end layered security

The WS5100's comprehensive security includes integrated features such as intrusion detection, an IPSec VPN gateway, AAA/RADIUS server (for WPA/WPA2 termination on the box) and "hotspot" provisioning capabilities for secure guest access. The stateful packet inspection firewall offers protection against denial of service attacks while optimizing network traffic. With support for the wireless security standards of today and the ability to easily upgrade to tomorrow's standards, the WS5100 delivers true value.

Simplified, centralized management

The WS5100 provides unified management of network hardware, software configuration, and network policies, and has built-in process monitors and troubleshooting tools. Motorola's Mobility Services Platform (sold separately) provides both device level management capabilities and centralized management of the WS5100 infrastructure in distributed locations. With active/active failover and clustering capabilities, as well as mobile unit load balancing, the WS5100 maximizes network uptime while minimizing network latency. Each WS5100 supports up to 48 access ports and 32 WLANs.

For more information, visit us on the web at www.motorola.com/ws5100 or access our global contact directory at www.motorola.com/enterprise/contactus

SPECIFICATION SHEET

WS5100

multi-protocol, 802.11n, 802.11g, 802.11b wireless switch

Clustering and load balancing

Ensures loads are balanced between access ports to ensure quality application performance, supports multiple levels of redundancy in case of failure

Mobility enablers

Virtual AP, Pre-emptive roaming, transmit power control, power save protocol, self-healing (triggered on loss or disruption of RF coverage). "Virtual AP" provides better control of broadcast traffic and enables multiple mobile and wireless applications with quality of service when network is congested. Pre-emptive roaming ensures Motorola mobile devices roam before signal quality degrades; PSP optimizes battery life, self-healing provides continuous network coverage in the event of disruption

Quality of Service (QoS)

Enhanced voice and video capabilities; prioritizes network traffic to minimize latency and provide optimal responsiveness to all users; Wi-Fi Multimedia Extensions (WMM-Power Save with Admission Control) for enhanced multimedia application support; improved battery life and capacity

WS5100 Specifications

Packet Forwarding

802.1D-1999 Ethernet bridging; 802.11-802.3 bridging; 802.1Q VLAN tagging & trunking; proxy ARP; IP packet steering-redirection

Wireless Networking

Wireless LAN:	Supports 32 WLANs; multi-ESS/BSSID traffic segmentation; VLAN to ESSID mapping; auto assignment of VLANs (on RADIUS authentication); power save protocol polling; pre-emptive roaming; congestion control with Bandwidth Management; VLAN Pooling
Access ports:	Supports 1-48 "thin" access ports; automatic access port adoption with ACLs; access port load balancing; direct sequence access point-to-access port conversion
Adaptive AP:	Supports 1-256 adoption of the Independent Motorola AP51X1 Access Point in Adaptive Mode for remote site and branch office solutions

Layer 2 or Layer 3 deployment of access ports

Layer 3 mobility (inter-subnet roaming)

Supported access ports and access points: Access ports – AP100 (802.11b) (L2 deployment only); AP300 (802.11a/b/g ready) (L2 or L3 deployments) with Static IP support

Access points – AP51X1 – Adaptive AP mode; AP4131 (L2 deployments only)

Radio frequency automatic channel select (ACS); transmit power control management (TPC); country code-based RF configuration; 802.11b – 3 non-overlapping channels; 802.11a-11 non-overlapping channels; 802.11g-3 non-overlapping channels

Network Security

Packet filtering/Access control lists (ACLs):	L2/3/4 stateful packet analysis; network address translation (NAT)
Authentication:	Access Control Lists (ACLs); pre-shared keys (PSK); 802.1x/EAP-transport layer security (TLS), tunneled transport layer security (TTLS), protected EAP (PEAP); Kerberos Integrated AAA/RADIUS server with native support for EAP-TTLS and EAP-PEAP (includes a built in user name/password database; supports LDAP)
Transport encryption:	WEP 40/128 (RC4), KeyGuard, WPA-TKIP, WPA2-CCMP (AES), WPA2-TKIP

IPSec VPN gateway (support for up to 100 tunnels): Supports DES, 3DES and AES encryption

Secure guest access (HotSpot provisioning): Local Web-based authentication; URL redirection for user login; customizable login/welcome pages; support for external authentication/billing systems

RADIUS support (standard and Motorola vendor specific attributes):

- User-based VLANs (standard)
- MAC-based authentication (standard)
- User-based QoS (Motorola VSA)
- Location-based authentication (Motorola VSA)
- Allowed ESSIDs (Symbol VSA)

NAC support with third party systems from Microsoft and Sygate

Optimized Wireless QoS

RF priority: 802.11 traffic prioritization and precedence

Wi-Fi multimedia extensions: WMM-power save with admission control

Classification & marking: Layer 1-4 packet classification; 802.1p VLAN priority, DiffServ/TOS

System Resiliency & Redundancy

Active: Standby, Active/Active and 1:Many redundancy with access port and MU load balancing; self healing (on detection of RF interference or loss of RF coverage)

Management

Command line interface (serial, telnet, SSH); secure Web-based GUI (SSL); SNMP v1/v2/v3; SNMP traps-40+ user configurable options; Syslog; TFTP Client; secure network time protocol (SNTP); text-based switch configuration files; DHCP (client/server/relay), switch auto-configuration and firmware updates with DHCP options; multiple user roles (for switch access); Syslog, MIBs (MIB-II, Etherstats, wireless switch specific monitoring and configuration)

Physical Characteristics

Form factor:	Standard 1RU
Dimensions:	1.73 in. H x 16.89 in. W x 15.93 in. D 43.9 mm H x 429 mm W x 404.6 mm D
Weight:	13.75 lbs./6.25 kg
Physical interfaces:	RS232 serial console port; 10/100/1000 Ethernet ports
MTBF:	>75,000 Hours

Power Requirements

AC input voltage:	100-240 VAC
Max AC input current:	6A@115 VAC, 3A@230 VAC
Max power consumption:	100-240 VAC, 50/60 Hz, 3A, 240 VAC, 50/60 Hz, 1.5A
Input frequency:	47 Hz to 63 Hz

User Environment

Operating temperature:	50°F to 95°F/10°C to 35°C
Storage temperature:	40°F to 149°F/-40°C to 65°C
Operating humidity:	5%-85% (w/o condensation)
Storage humidity:	5%-95% (w/o condensation)
Operating altitude:	50 ft. to 10,000 ft./16 m to 3,048 m
Storage altitude:	50 ft. to 35,000 ft./16 m to 10,600 m

Regulatory

Safety certifications:	FCC (Art.15, part B), Industry Canada, CE, VCCI, C-Tick, BSMI
EMI compliance:	UL 1950, cUL (Canada), VDE GS, DENAN (Japan), CB Cert

Part Numbers

WS-5100-06-WWR	6 Port WS5100 Wireless Switch
WS-5100-12-WWR	12 Port WS5100 Wireless Switch
WS-5100-18-WWR	18 Port WS5100 Wireless Switch
WS-5100-24-WWR	24 Port WS5100 Wireless Switch
WS-5100-30-WWR	30 Port WS5100 Wireless Switch
WS-5100-36-WWR	36 Port WS5100 Wireless Switch
WS-5100-42-WWR	42 Port WS5100 Wireless Switch
WS-5100-48-WWR	48 Port WS5100 Wireless Switch
WS-5100-RS-WWR	Redundant WS5100 Wireless Switch
WS-5100-UC-WW	6 Port Upgrade

**MOTOROLA**

motorola.com



Part number SS-WS5100. Printed in USA 02/08. MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. ©2007 Motorola, Inc. All rights reserved. For system, product or services availability and specific information within your country, please contact your local Motorola office or Business Partner. Specifications are subject to change without notice.

EXHIBIT 19



DATA SHEET Aruba Mobility Management System

ARUBA MOBILITY MANAGEMENT SYSTEM

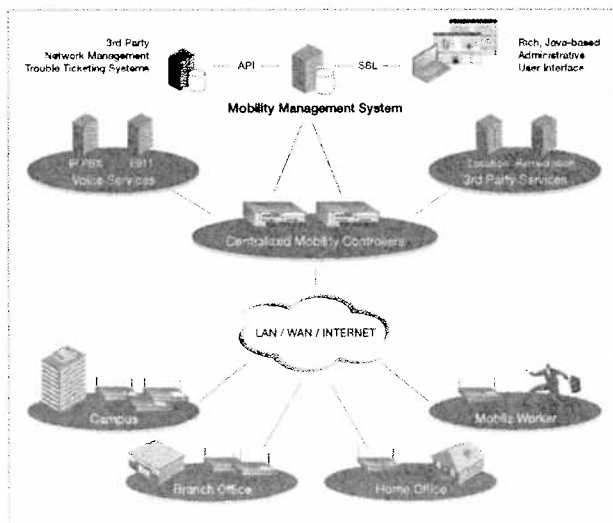
The Aruba Mobility Management System provides a comprehensive suite of applications for planning, configuration, fault and performance management, reporting, RF visualization and Wi-Fi device location identification for Aruba's user-centric networks. It seamlessly integrates with Aruba's access points and Mobility Controllers and supports the new paradigm of adaptive wireless LANs, identity-based security, and application continuity, in contrast with previous generations of port-based network management. The user-centric management model allows IT administrators to rapidly visualize all network objects related to the user in real time. This approach drastically cuts down the mean time to resolution (MTTR) and ensures a high-quality WLAN user experience.

The Mobility Management System reduces total cost of ownership by automatically discovering and managing hundreds of controllers and thousands of access points (APs) and users from a single network operations center. Centralized configuration management, coupled with the ability to track client devices, identify rogue devices, and plan new deployments and visualize RF coverage patterns with an intuitive, seamless UI, is a key differentiator. The Mobility Management System comes with a built-in location API that enables external systems to query the location of any WLAN device. The Mobility Management System software can be deployed on any PC platform (Linux or Windows 2003/XP) or optionally purchased as an enterprise class, hardened appliance.

Aruba's Mobility Management System consists of six integrated Java-based client-server applications, namely: Dashboard, Monitor, Reports, Configuration, RF Live, and RF Locate. Additionally, the Java client applications are auto-updated using the Java Web Start capability which eliminates the need to manage client desktops and laptops independently.

From the Dashboard, the administrator can launch monitoring, reporting or RF views for any resource and jump directly into context sensitive detailed views. This streamlines the problem isolation and resolution effort. Information such as user counts, controller modes, AP types/mode (Rogue, Valid, Interfering) can be viewed within the Mobility Management System or can be exported for use in custom applications.

DASHBOARD



The Dashboard application gives network administrators a summary view of the network infrastructure. Minor, Major and critical alarms for fault management are displayed for every physical or logical resource on the network. The network administrator can navigate to any resource by browsing through a tree view and instantaneously assess the health of the WLAN infrastructure. With a single mouse click and using a robust flexible search and filter engine, a network administrator can rapidly zero in on a specific client or AP and determine its attributes.



Aruba MMS Dashboard

MONITOR

The Monitor application provides real-time polling and visualization of Aruba Mobility Controller and AP data via secure SNMPv3. The live graphing capability of the Monitor application allows the network administrator to select any object and create a live graph for object attributes. Data can be exported for custom processing or use in third-party applications.

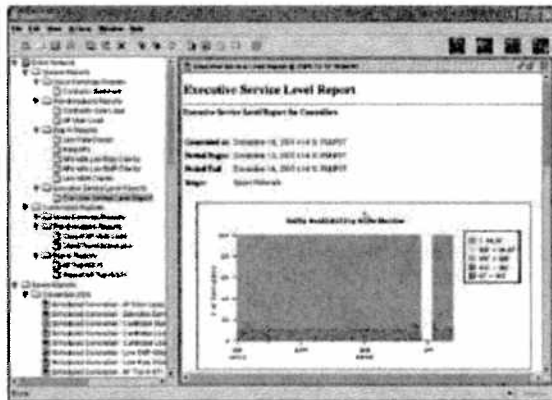
ARUBA MOBILITY MANAGEMENT SYSTEM



Aruba MMS Monitor

REPORTS

The Reports application allows network administrators to configure, schedule and run reports for custom time intervals. Historical trend reports dating back several weeks and months can be generated by leveraging WLAN usage and performance data, continually collected and stored within the Mobility Management System. These reports can be reviewed by IT managers interested in WLAN, network health, performance, usage and capacity planning. The flexible reporting framework supports pre-defined TopN, Trend and Summary reports, in addition to user configurable custom reports. The reports can be scheduled to run at periodic intervals and automatically e-mailed. The report output is in industry standard HTML and PDF formats.



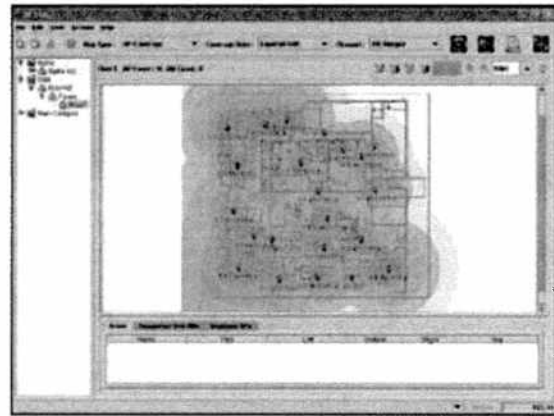
Aruba MMS Reports

CONFIGURATION

The Configuration application supports options to define user- and role-based access control policies. The Aruba WLAN infrastructure can be grouped into hierarchical "equipment contexts," "managed domains," and "controller" and "access point" groups. Configuration parameters are conveniently categorized into profiles. Profiles can be configured independently for WLAN authentication, access control, SSID, RF management, intrusion prevention, firewall policies, IP mobility, captive portal and other parameters. Controller and AP groups can be configured to reference profile objects. Profiles applicable at a global level can be separated from custom profiles applicable to equipment deployed at specific geographic locations.

RF LIVE

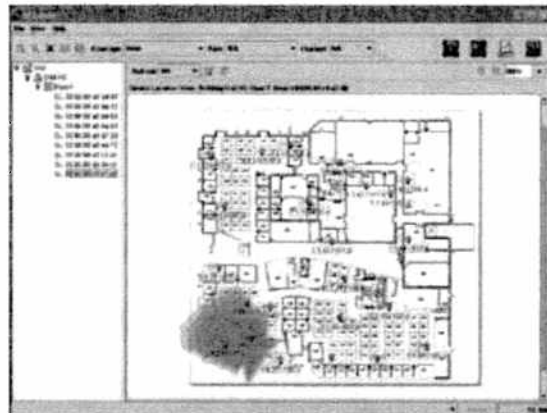
The RF Live application is used both for pre-deployment planning and live visualization of the RF environment and is useful for analyzing interference patterns and coverage holes. RF parameters such as signal strength and interference are displayed in the context of floor plans with superimposed coverage contours and colored heat maps. Since RF Live uses dynamic information delivered by Aruba's patent-pending Adaptive Radio Management™ (ARM) algorithms, it provides a real-time understanding of the evolving RF environment and eliminates the need for manual, post-deployment RF fingerprinting. RF Live also provides integrated tools for pre-deployment RF planning. AutoCAD drawing (.dwg files) import allows network administrators to easily load existing floor plans into the system to facilitate the planning process. Imported floor plans can be used to determine ideal access point placement based on coverage and capacity requirements. The planning process is straightforward since RF tuning is a dynamic, real-time process managed by the controller. ARM capability eliminates the need for traditional heavy duty planning tools that require a detailed understanding of building materials or an expensive manual site survey.



Aruba MMS RF Live

RF LOCATE

The RF Locate application can track and locate any Wi-Fi device within range of the Aruba mobility infrastructure. Using accurate deployment layouts and triangulation algorithms, the RF Locate application enables the network administrator to rapidly locate selected client devices. Devices that can be easily located include PDAs, rogue APs/Clients, VoWLAN phones, laptops, and Wi-Fi asset management tags. Aruba offers an open XML/SOAP web services API to allow third-party applications to access location data enabling sophisticated next-generation location-based applications.



Aruba MMS RF Locate

ARUBA MOBILITY MANAGEMENT SYSTEM

FEATURE	BENEFIT
User-centric Data Model	Rapid problem isolation and resolution
Auto-discovery	Auto-discovery based on an IP address seed
User configurable reports	Flexibility to create and visualize WLAN usage, availability and performance data by controllers and APs
Versatile scheduling engine	Run reports, locate stations, email reports at scheduled intervals
Powerful visualization	Integrated fault and performance management, RF coverage and interference visualization
Location API	Enables 3rd party applications to support innovative location-based services
Saved searches	Allows network administrators to save and re-use frequently used search patterns. Accelerates mean time to resolution
Java™ Webstart UI	On demand UI download provides "client anywhere" flexibility with performance advantages of a thick client

ORDERING INFORMATION

The Aruba Mobility Management System is available both as an integrated appliance and as a software application. The MM-100 Mobility Management System Appliance is a high-performance system with pre-loaded software and support for 250 APs. The software version ships with support for 50 APs. Both versions can be upgraded in increments of 100 APs. Unlimited AP Expansion licenses are also available for both platforms.

PART NUMBERS	PRODUCT DESCRIPTIONS
MM-100	Aruba MM-100 Mobility Management System Appliance (Up to 250 APs) Contains: MM-100 Appliance (Dual-processor, RAID subsystem, dual-power supply, 2GB RAM, 1U), 1 rackmount kit, 1 Ethernet cable, software restore disk, quick start installation guide, user guide, front bezel, 2 power cords. Includes license activation certificate for 250 APs.
MM-SW	Aruba Mobility Management System Software (Up to 50 APs) Contains: Software distribution disk, quick start installation guide, user guide. Management of up to 50 concurrent APs. Includes license activation certificate for 50 APs. Requires Pentium 4 3.0GHZ 800FSB, 2GB DDRII RAM, 200GB SATA/SCSI HD, RedHat ES 4.0 Hardware Compatibility Compliance. Dedicated system required.
MM-SW-WIN	Aruba Mobility Management System Software for Windows (Up to 50 APs) Contains: Software distribution disk, quick start installation guide, user guide. Management of up to 50 concurrent AP. Includes license activation certificate for 50 APs. Requires Pentium 4 3.0GHZ 800FSB, 2GB DDRII RAM, 200GB SATA/SCSI HD. Dedicated system required.
LIC-MM-100	Mobility Management System Expansion License - Add 100 AP Each LIC-MM-100 upgrade adds management support for 100 additional APs. Upgrades are cumulative and multiple licenses can be applied to the same system. Managed APs can be local or remote.
LIC-MM-UL-1	Unlimited AP Expansion License for MM-SW Mobility Management System Software The LIC-MM-UL-1 upgrade provides unlimited AP support on a single MMS-SW Mobility Management System. Managed APs can be local or remote.
LIC-MM-UL-2	Unlimited AP Expansion License for MM-100 Mobility Management System Appliance The LIC-MM-UL-2 upgrade provides unlimited AP support on a single Aruba MM-100 Mobility Management System Appliance. Managed APs can be local or remote.



WWW.ARUBANETWORKS.COM

1322 Crossman Avenue, Sunnyvale, CA 94089 Tel. +1 408.227.4500 Fax. +1 408.227.4550

EXHIBIT 20



RF Design and Management Software: Motorola EnterprisePlanner™ Plan and Optimize Multiple Wireless Networks




The large number of wireless frequencies, environmental factors, and network equipment choices presents daunting challenges to even the most experienced wireless network designer. Building a wireless network that delivers exceptional Quality of Service (QoS) and maximum value means designing it for the context in which it operates and the capacity and coverage that your business requires.

Motorola EnterprisePlanner™ enables rapid design of customized wireless networks with the capacity, reliability, and performance required to run the most demanding applications, even in complex RF environments. EnterprisePlanner maximizes wireless network performance, ensuring that your critical wireless applications are effectively deployed and fully functional.


At the heart of EnterprisePlanner is an RF-intelligent, site-specific 2D/3D model of your facility. Motorola's patented predictive software uses this model to accurately simulate your wireless network's coverage and capacity. Using a library of data amassed through a decade of research and thousands of implementations, EnterprisePlanner allows you to graphically visualize the impact of building materials, the intended use of the network, and the sources of radio frequency (RF) interference. As a result, EnterprisePlanner can dramatically reduce wireless network planning, deployment, and operating costs.

Validating Wireless Coverage, Boosting Performance


EnterprisePlanner combines advanced facility modeling and predictive algorithms with intuitive design wizards, providing you the power to efficiently design and simulate a wide range of multi-band indoor and outdoor campus networks including WLAN, 3G, CDMA, GSM, iDEN, WiMAX and user defined frequencies. EnterprisePlanner allows you to import building and site information from a variety of formats including AutoCAD® files, scanned images, and GIS maps. This model becomes RF-intelligent once you map the RF characteristics of interior and exterior walls and other potential obstructions to Motorola's extensive RF attenuation database.



Quickly and easily create 2D/3D building models from CAD drawing files, scanned blueprints or rapid sketching



Simulate a mobile client in a GSM network



Run "what if" scenarios to visualize WLAN coverage

Enabling Application-Specific Wireless Network Design

With EnterprisePlanner's RF-intelligent model in place, you can then establish user regions to designate the expected network traffic and the types of applications in use. Each user region can be set up for multiple functions. For example, an area might have 50 low-bandwidth email users, 20 medium-bandwidth users performing internal file transfers, and 10 high-bandwidth voice over WiFi (VoWi-Fi) users. EnterprisePlanner considers all of these factors — environment, users, and applications — and creates a network plan that includes the number of access points, equipment placement, channel settings, and power levels required to support the usage requirements.

With EnterprisePlanner, you can prevent wireless coverage gaps, mitigate the effect of signal interference, and ensure maximum capacity without the need for specialized in-house RF expertise.

Simplifying Design,

Maximizing Return on Investment

Wireless networks are not static, and performance can fluctuate over time. Motorola's predictive design approach eliminates this challenge by empowering you to optimize the network before spending a single dollar on hardware or wasting time on trial-and-error deployments. Planning teams can quickly access equipment location data, configuration information, and run "what if" scenarios to simulate potential network changes. Network deployment plans are easily stored in a centralized database, making network enhancements and expansions simple to evaluate and easy to plan.

Features and Benefits:

- Tailor your network for a wide range of wireless standards including WLAN, 3G, CDMA, GSM, iDEN, WiMAX, and user defined frequencies
- Directly import AutoCAD, PDF, JPEG, BMP and any other common building or site map file format
- Create an RF-intelligent model by leveraging an extensive RF attenuation database and network parts list library
- Automated access point/base station placement and configuration
- Visualize the physical location and configuration of all network equipment
- Review QoS critical information such as RSSI (Received Signal Strength Intensity), SIR (Signal to Interference Ratio), SNR (Signal to Noise Ratio), and throughput
- Predict how radio frequency activity will impact your wireless network
- Optimize your wireless network by leveraging site-specific wireless phone and RF receiver activity
- Automatically generate bill-of-materials and maintenance records for use by deployment teams and in future network expansion
- Reduce the total cost of ownership of wireless networks by eliminating costly rework

The Motorola Advantage

Motorola pioneered voice communication and mobile communication technology. This experience enables us to help you own, design, customize and control your own network—thereby increasing efficiency, interoperability and security throughout your entire enterprise.



Motorola, Inc. 1301 East Algonquin Road, Schaumburg, Illinois 60196 U.S.A.
1-800-367-2346 motorola.com/enterprise
RC-10-2016

MOTOROLA and the Stylized M Logo are registered in the U.S. Patent and Trademark Office. AutoDesk is a registered trademark of AutoDesk, Inc. All other product or service names are the property of their respective owners. © Copyright Motorola, Inc. 2006

EXHIBIT 21

RF Design and Management Software: Motorola LANPlanner®

Ensuring High Performance WLAN Networks



Building a wireless network that delivers superior Quality of Service (QoS) and maximum value from your investment means designing it for the conditions in which it operates and the capacity and coverage that customers require. The multitude of wireless applications, impact of environmental factors, and variety of network equipment configurations can present daunting challenges to even the most experienced wireless network designer.

Predictive Design for WLAN Networks
Motorola LANPlanner® enables network planners to rapidly and accurately design robust wireless networks with the capacity, reliability, and performance required for implementing the most demanding wireless applications, in the most challenging environments. LANPlanner takes wireless network performance to the next level, ensuring that critical wireless applications and infrastructure are effectively deployed, fully functional.

At the heart of LANPlanner is an RF-intelligent, site-specific 2D/3D model of a deployment facility. LANPlanner allows network designers to predict and visualize the impact of construction materials, the intended use of the network, and the potential impact of co-channel interference. As a result, LANPlanner can reduce wireless network planning, deployment, and operating costs—while providing superior wireless QoS.


RF-Intelligence Enables Optimal Wireless Network Performance
LANPlanner's easy to use GUI allows you to quickly import building and site information from a variety of formats including AutoCAD® files, scanned images, and hand drawn floor plans. This import becomes RF-intelligent after users map the characteristics of interior and exterior walls and other potential wireless obstructions to Motorola's RF attenuation library.

Fast and Efficient WLAN Planning for Demanding Applications
After creating the RF-intelligent model, designers add information about the number of users, their locations, and the applications they plan to use. LANPlanner then combines this information with the RF-intelligent model to recommend the number and placement of Access Points (APs).

"An investment in network design can affect a 50% savings in WLAN setup costs."

—Gartner

"Wireless LAN Technology Scenario"



Run "what if" scenarios to visualize WLAN coverage



Survey the deployed network to validate performance

LANPlanner Reporting and Measurement Capabilities Save Time

Access point configuration can be difficult, especially for large wireless network deployments. To save time and money during deployment, LANPlanner generates reports in Microsoft Word® format that contain access point locations, data rate requirements, channel assignments, power levels, and SSID allocations. After deployment, teams can also collect network performance data and visualize heat maps of the information they collected, validating that the deployed network performs to expectations.





From Requirements Gathering to Post Deployment Verification

LANPlanner is a key driver of Motorola's Wireless Network Design Process. With this software, users are able to perform Phase 1, 2 and 3. This approach ensures that network planning teams consider the impact of applications, user density, and the deployment environment on network performance. The design plan can also be used to simulate network expansion or confirm the viability of adding new wireless applications.

Features and Benefits

- Rapidly define user throughput regions and automatically place access points and sensors to achieve a network design plan with user-driven coverage and capacity requirements
- Configure access points and sensors and predict how RF activity will impact your wireless network performance and QoS
- Graphically visualize the physical location and configuration of all installed network equipment
- Review QoS critical information such as RSSI (Received Signal Strength Intensity), SIR (Signal to Interference Ratio), SNR (Signal to Noise Ratio), and user region data rates
- Verify post-deployment network performance and visualize heat maps of measured data
- Quickly load AutoCAD, PDF, JPEG, and any common building or site map file format and create a reusable, extensible RF-intelligent model
- Automatically generate bill-of-materials and maintenance records for use by deployment teams and in future network expansion
- Reduce the total cost of ownership (TCO) of wireless networks by eliminating the costly rework that frequently occurs with measurement based network design.

Wireless Network Design Process

PHASE 1	PHASE 2	PHASE 3	PHASE 4
 <p>Requirements Gathering</p> <ul style="list-style-type: none"> • Secure AutoCAD/Facility Diagrams (indoor networks) • Gather Network Requirements: Number of Users, Applications, and Equipment Preferences 	 <p>Design</p> <ul style="list-style-type: none"> • Create RF-Intelligent Software Model • Plan for Coverage and Capacity Requirements • Place Wireless Hardware • Print Network Design Plans • Generate Bill of Materials 	 <p>Deployment</p> <ul style="list-style-type: none"> • Deliver Deployment Plan • Advise Deployment Team (in-house or third-party) 	 <p>Verification and Optimization</p> <ul style="list-style-type: none"> • Measure Coverage • Optimize RF-Intelligent Model • Fine-tune Wireless Network for QoS • Verify Performance Specifications are Met

The Motorola Advantage

Motorola pioneered voice communication and mobile communication technology. This experience enables us to help you design, customize, and control your own network—thereby increasing efficiency, interoperability, and security throughout your entire organization.



Motorola, Inc. • 1301 E. Algonquin Road • Schaumburg, Illinois 60196 U.S.A.
www.motorola.com/rfdesign • 1-800-901-6484 • 1-512-821-1560

RC-10-2021

MOTOROLA and the Stylized M Logo are registered in the U.S. Patent and Trademark Office. All other product or service names are the property of their respective owners. © Copyright Motorola, Inc. 2006

EXHIBIT 22

- o Careers
- o Contact Us

September 10, 2007

[Company](#) > [News & Events](#) > September 10, 2007

ARUBA GAINS MARKET SHARE AND SOLIDIFIES POSITION AS THE WORLD'S SECOND LARGEST ENTERPRISE WIRELESS LAN SUPPLIER

Sunnyvale, CA, September 10, 2007 - Aruba Networks, Inc. (NASDAQ: ARUN), a global leader in user-centric networks and secure mobility solutions, today announced that its share of the enterprise wireless LAN market has risen to greater than 10% in the second quarter of 2007 from roughly 5% in the same period of 2005. The market data are based on a newly published report by Dell'Oro Group, and include sales of Aruba products by Alcatel-Lucent. During the same period Motorola's Symbol unit lost market share, and Aruba displaced Motorola as the world's second largest enterprise wireless LAN supplier.

Aruba's user-centric networks integrate adaptive wireless LANs, identity-based security, and application continuity services into a cohesive, high-performance system that securely delivers the enterprise network to users, wherever they work or roam. By extending the enterprise to reach all users without compromising security or convenience, Aruba has redefined mobility with respect to where and how people work.

"Aruba has been a leading innovator in the wireless LAN market, and we have been rewarded with strong growth in our customer base, revenue, and market share," said Keerti Melkote, Aruba's co-founder and head of products and partnerships. "It is clear that the enterprise wireless LAN market is becoming a two horse race as we continue to gain market share at the expense of incumbents and struggling small suppliers."

With regard to the lawsuit filed by Motorola on the eve of Aruba's end-of-quarter results, during which the company announced a significant increase in revenue, Melkote stated, "We do not believe that we infringe Motorola's patents, and we intend to vigorously defend ourselves against the claims when we have our day in court. We also look forward to continuing to compete successfully against Motorola in the market."

Aruba is a member of a coalition including Cisco, Intel, Microsoft, SAP, Oracle, eBay, and Symantec, among others, supporting The Patent Reform Act of 2007, which passed the House on September 7 by a vote of 220-175. The Act aims to overhaul the U.S. patent system, which the coalition believes is overburdened with applications for trivial inventions and patent litigation intended to prevent competition and reap windfall profits.

About Aruba Networks

Aruba securely delivers the enterprise network to users, wherever they work or roam, with user-centric networks that significantly expand the reach of traditional port-centric networks. User-centric networks integrate adaptive WLANs, identity-based security, and application continuity services into a cohesive, high-performance system that can be easily deployed as an overlay on top of existing network infrastructure. Adaptive WLANs deliver high-performance, follow-me connectivity so users are always within reach of mission-critical information. Identity-based security associates access policies with users, not ports, to enable follow-me security that is enforced regardless of access method or location. Application continuity services enable follow-me applications that can be seamlessly accessed across WLAN and cellular networks. The cost, convenience, and security benefits of user-centric networks are fundamentally changing how and where we work. Listed on the NASDAQ and Russell 2000 Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at <http://www.arubanetworks.com>.

This press release may contain statements relating to future plans, events or performance. Such statements may involve risks and uncertainties, including risks associated with uncertainties pertaining to obtaining approvals for and satisfying conditions to the consummation of Aruba Network's acquisition of Network Chemistry's assets, the actual impact of the acquisition on operating results, including the purchase price allocations for assets and liabilities, the amount of IPR&D and other intangible assets, and the time period in which intangible assets will be amortized, the successful integration and support of Network Chemistry's technology and employees into Aruba, Aruba's ability to incorporate the Network Chemistry technology into Aruba's mobility solutions in a timely and cost-effective manner, market acceptance of the RFprotect and BlueScanner products, the ability of the RFprotect and BlueScanner products to function as designed, the ability of Aruba to support RFprotect and BlueScanner products, the timing and level of customer orders, demand for products and services, the development of markets for Aruba's products and services, and other risks identified in Aruba's SEC filings. Actual results, effect on earnings, events and performance may differ materially. Readers are cautioned not to place undue reliance on these forward-looking statements, which speak only as of the date hereof. Aruba undertakes no obligation to release publicly the result of any revisions to these forward-looking statements that may be made to reflect events or circumstances after the date hereof or to reflect the occurrence of unanticipated events.

###

Media Contacts

Michael Tennefoss
Aruba Networks, Inc.
+1-408-754-8034
mtennefoss@arubanetworks.com

Patty Oien
Breakaway Communications.
+1-415-358-2482
poien@breakawaycom.com

© 2008 Aruba Networks, Inc. All rights reserved. Specifications are subject to change without notice.

Archive

- **Latest Press Releases**
- **2008**
- 2007
- 2006
- 2005
- 2004
- 2003

- Support
- Site Map
- Privacy Policy
- Terms of Use
- Legal

Copyright © 2008. Aruba Networks, Inc. All rights reserved.

EXHIBIT 23


This page should print automatically. If not select **Print** from the **File** menu.

Enterprise Wi-Fi - the market shifts again

Jockeying for position as we prepare for a boom

Peter Judge, Techworld September 21, 07

Advertisement



The enterprise Wi-Fi market is changing gear again. Vendors selling Wi-Fi networks have scented a big expansion opportunity driven by 802.11n. The fast Wi-Fi standard - still not complete, but very much a marketing reality - will finally make Wi-Fi a real competitor to Ethernet in connections to end-users' office desktops and laptops.

So the vendors all sat up and took more notice than usual of this quarter's market share pronouncements. After all, their position this quarter could determine the share they get of the big 802.11n Wi-Fi boom.

"If you were an IT manager thinking which company you might go with, you might be less inclined to invest in a company with a smaller share," said Elmer Choy, Wi-Fi market share analyst at Dell'Oro whose figures have sparked a flurry of releases. "It could influence buying decisions."

Who's number two?

As always, there's no doubt who's top. Cisco has the biggest revenue from enterprise Wi-Fi equipment by far. It's got about 60 percent of the enterprise Wi-Fi market, according to figures from Dell'Oro.

A lot of this - maybe even the majority - will still be from old-style standalone Aironet access points, not Wi-Fi switches. However, the giant is actually still gaining share. It only had around 50 percent in 2002, but bought into Wi-Fi switches in time to keep up with the trends. It's recent announcement of an 802.11n access point is clearly designed to show it's on top of the technology.

Despite Cisco's lead, Motorola surprised us with a breezy release claiming its Symbol subsidiary is actually number one in wireless switches. Now Symbol invented the wireless switch, but it's definitely not the leading Wi-Fi player overall. Motorola's claim only applies to switches, not APs, and it turns out to refer to market share by unit. It sold 33 percent of the Wi-Fi switches in the last quarter (and also, by coincidence, sold its 100,000th switch).

A lot of those switches will be the WS2000 switch, a nice branch product which can power four access points (and manage two more with PoE injectors). "Motorola is still shipping to their strength in the retail vertical market, usually with a handful of APs per switch," said Choy. "That's very much their market."

The WS2000 only costs \$1000 and goes a long way to explaining how Motorola is shipping so many switches, and still lying third in revenue.

Aruba is chuffed

We're giving a bit more credence to Aruba's claim that it overtook Symbol in market share by revenue, and is number two in the Wi-Fi market. In perspective, Choy reckons it has about ten percent of the market to Symbol's eight percent. To get this figure, Aruba added in the OEM revenue it gets from products sold through Alcatel (and some other ones sold under badge deals with the likes of Netgear).

The figures are backed up by Infonetics Research: "Aruba is chuffed at being second," said Infonetics' wireless analyst Richard Webb. "We've had them neck and neck for a year now."

It's especially good news for Aruba to overtake a Motorola-owned Symbol, said Webb: "I thought Symbol sales would drive forward, but they still don't have great reach into enterprise markets."

Aruba's obviously also keen to get one over on Motorola, in the light of Motorola's allegation that Aruba infringed its patents.

Was Motorola's OEM business also included in Dell'Oro's figures? Apparently it was, even though Motorola/Symbol has always been strangely reticent about this, but it has deals with IBM and HP.

More contenders

Further down the list there's another surprise vendor shouting. Meru, a vendor with a "blanket" Wi-Fi architecture claims to be in fourth place. That's not something a vendor would normally get excited about, but if it really is there, an upstart with a new architecture appears to have barged past an established vendor most analysts usually place third in the enterprise Wi-Fi market - Trapeze.

"It's the first time Meru has reported figures," says Choy. "It's just ahead of a pack of smaller players." That pack includes 3Com, Alcatel Lucent, Enterasys, Extreme, HP Procurve, Proxim, Bluesocket and others. The figures are too uncertain to quote for these ones.

Webb also praises Meru, as an "innovative voice in the WLAN market." It's clearly had a good time - it's claiming to have doubled its revenue between the first and second quarter of 2006, but I'm dubious that it's really ahead of Trapeze.

That has to be a guess, because Trapeze has never quoted figures to analysts (and has been promising to do so "soon" for about three years). But Trapeze has been established for more than five years, and gets about half its revenue from OEM deals with the likes of Nortel, 3Com and Enterasys, each of which apparently does too little Wi-Fi business to stand out from the pack.

"Anecdotally, Trapeze is doing very well," concedes Choy. Webb agrees, but he's sure Trapeze has been left behind by an Aruba buoyed up by its recent stock market flotation.

Trapeze for sale?

Just how well Trapeze is doing could be important, if there's any truth in the rumours that Trapeze is up for sale. The rumour comes from wireless site Unstrung, which suspects Trapeze's partners 3Com and Nortel (despite the Canadian giant signalling its intent to end its OEM agreement with Trapeze), as well as Juniper, which has invested in Trapeze before.

What about the others?

There are still other players. One worth remembering is Colubris. It doesn't rate much excitement in Dell'Oro's figures, but Choy concedes it may be a bit under-reported, as some of its products go to service providers, and might not show up in an enterprise share study. "It's definitely growing," he said.

And there are still new players yet to register. Aerohive, a company which takes the radical step of doing away with the central controller in favour of a self-organising set of co-operative nodes, received funding in the summer and is expanding its sales effort - we expect to hear more from them as well.

With an 802.11n-driven boom in the offing, it's not going to be quiet in the WLAN world.

This article was printed from **Techworld** : www.techworld.com The UK's infrastructure & network
knowledge centre © 2006 : All rights reserved

[Click here to close this window and return to the website](#)

EXHIBIT 24

Seeking Alpha^α

IPO Analysis: Aruba Networks Doesn't Have Enough Under The Hood

by: Bill Simpson

posted on: April 03, 2007 | about stocks: [ARUN](#)



On March 20, [Bill Simpson](#) wrote an analysis of Aruba Networks ([ARUN](#)). On March 26, the company raised \$88 million with an offering of 8 million shares at \$11, above the forecasted range of \$8 to \$10. On Tuesday April 3, ARUN closed at \$14.10.

The text of Mr. Simpson's original writeup follows:

...

Aruba Networks plans on offering 9.2 million shares (assuming over-allotments) at a range of \$8-\$10. Goldman Sachs and Lehman are lead managing, JP Morgan and RBC Capital co-managing. Post-offering ARUN will have 76.4 million shares outstanding for a market cap of \$688 million on a \$9 pricing.

Note - In a trend we've seen with a few ipos lately, ARUN has excessive option/warrant/employee incentive shares already outstanding and ready to be converted to shares. ARUN has 19.9 million options outstanding with an average exercise price of \$2.71 per share. In addition, ARUN has approximately 1.6 million warrants and employee incentive shares already awarded. These 21.5 million options/warrants/employee incentive awards will be converted into shares sooner than later and must be added into the initial market cap. Factoring in this upcoming dilution, ARUN will have 97.9 million shares for a market cap of \$881 million on a pricing of \$9.

IPO proceeds will be used for working capital and general corporate purposes.

4 venture capital firms will own between 55%-60% of ARUN post-ipo. Fully expect these firms to divest a portion of their holdings at the 180 day lock-up expiry.

From the prospectus:

We provide an enterprise mobility solution that enables secure access to data, voice and video applications across wireless and wireline enterprise networks. Our Aruba Mobile Edge Architecture allows end-users to roam to different locations within an enterprise campus or office building while maintaining secure and consistent access to all of their network resources.

Yet another networking ipo whose primary competition is Cisco ([CSCO](#)) and Motorola ([MOT](#)). As we'll see in the financials, Cisco's size and ability to compete on price are greatly effecting ARUN's ability to book a net profit even while growing revenues.

ARUN's focus is enterprise WLAN or 'wireless local networks.' ARUN, which began shipping product in 2003, calls its product Aruba Mobile Edge Architecture [AMEA].

Industry - Enterprise wireless networking has grown appreciably due to the desire for mobile computing. WLAN or VPN (Virtual Private Networks) have been the solution enabling open access on wired network ports. These type networks extend the fixed network over the air. ARUN believes its AMEA product takes WLANs to another level by providing additional security features, allowing secure roaming over the entire network, increased performance, easy scalability and integration etc. Essentially ARUN believes it's built a better enterprise WLAN product.

According to ARUN:

Our architecture allows end-users to roam to different locations within an enterprise campus or office building while maintaining secure and consistent access to all of their network resources. Our architecture also enables IT managers to establish and enforce policies that control network access and prioritize application delivery based on an end-user's organizational role and authorization level.

ARUN's differentiators from traditional WLANs:

- 1 - Secure mobility - Enables secure roaming over the network in various remote locations.
- 2 - Improved application performance - non-fixed port user centric and application aware allowing prioritization and optimization of data, voice and video services based on the specific user and/or the application being delivered.
- 3 - Ease of deployment and integration - Designed as an overlay to existing enterprise networks, allowing customers to deploy ARUN's products without causing any disruption to existing network operations.
- 4 - Scalability - Can be scaled to support up to 100,000 concurrent users from a centralized point-of-control.
- 5 - Flexibility - Designed for ease in introducing new applications.

The majority of ARUN's revenues are derived through resellers, distributors and OEMs, not ARUN's direct sales force. Third parties accounted for 75% of 2006 revenues. Largest channel partner/re-seller of ARUN's products is Alcatel-Lucent (ALU) which accounted for 18% of revenues the past 18 months. ARUN has two interesting agreements with Alcatel-Lucent:

- 1) ARUN is restricted from selling products to other 3rd party sellers without consent of Alcatel-Lucent. Lot of power for a company that was responsible for just 18% of revenues in the past 18 months.
- 2) A 'most-favored nations' clause in which Alcatel-Lucent is guaranteed a match of the lowest price-point ARUN is selling its product to another channel partner/re-seller.

ARUN also has a strategic relationship with Microsoft, which began in June 2005. Microsoft chose ARUN's products for a worldwide company deployment in Asia, North America and Europe. Sales to Microsoft have totaled \$3.5 million. In addition ARUN will essentially 'gift' MSFT approximately 400,000 shares of ARUN stock on ipo. Essentially MSFT is getting either free equipment or free ARUN shares from ARUN for choosing ARUN products.

ARUN's products have been sold to over 200 end customers. Flextronics (FLEX) handles the majority of ARUN's manufacturing.

Financials

\$1 per share in cash post-ipo, no debt.

Revenues have been growing steadily if not rapidly. For the past 5 quarters ending 1/31/07, ARUN's revenues were (in millions): \$20.1, \$20.1, \$24, \$24.5, \$26.6. Operating expenses have also grown steadily in the past 5 quarters. In fact, ARUN's operating expense growth has kept pace dollar for dollar pretty much with revenue growth. This really isn't what you want to see in a young growing company.

The past 5 quarters of ARUN's operating expenses ending with the 1/31/07 quarter (in million): \$12.4, \$12.6, \$14.7, \$14.8, \$16.4. the % of operating expenses to revenues the past 5 quarters are 62%, 63%, 61%, 60%, 62%. 5 quarters of solid revenue growth, yet ARUN is no closer to lowering its operating expense ratio.

ARUN's fiscal year ends 7/31 annually. FY '07 then will end 7/31/07. Through the first 6 months, ARUN's revenues are on pace for approximately \$105-\$110 million. This would represent a 48% increase over FY '06. Keep in mind ARUN was essentially in start-up revenue stage as recently as FY '04. Also at that revenue run rate, ARUN will be selling a hefty 9 X's 2007 estimated revenues.

Gross margins appear on track for the 60% ballpark. As noted above, operating expense levels are where ARUN runs into trouble. By growing its operating expenses in the ballpark of dollar for dollar, ARUN is not shifting close to profitability even with the increased revenue growth. This is a very competitive niche and it appears ARUN is plowing a substantial amount of expense money into sales and markets to compete with Cisco. The result is that ARUN will actually lose more money in FY '07 than in FY '06 even with the 44% revenue growth.

Assuming the fully diluted 98 million shares outstanding, ARUN lost \$0.12 in FY '06 (ending 7/31/06) and appear on pace to lose \$0.20-\$0.25 in FY '07. Until ARUN is able to reduce its operating expense ratio significantly, it will not be able to turn a profit.

Thus far ARUN has not been able to lower that ratio much at all.

Conclusion - ARUN appears more than fully valued when all those option/warrant/employee incentive shares are factored in. These are shares that will be exercised over the next 1-2 years, so you've got to include them in the initial market cap. ARUN to me does not look anything like a \$1 billion market cap company should look. We've seen a number of good tech ipos the past 6 months, ARUN at this initial market cap is not one of them. There just isn't remotely enough here under the hood to justify a a nearly \$1 billion dollar cap on pricing.

EXHIBIT 25

March 26, 2008



Americas: Communications Technology

Transitioning CommTech coverage; lowering view to Neutral

Changes in coverage and ratings

We are transitioning coverage of the Communications Technology sector and lowering our coverage view to Neutral from Attractive. Simona Jankowski is assuming coverage of Cisco (adding Buy-rated CSCO to the Conviction List), Juniper (downgrade to Neutral), and Acme Packet (Neutral), and will assume coverage of Motorola (Not Rated), QUALCOMM (Buy), and Research In Motion (Buy) after interim coverage by Thomas Lee; Thomas Lee is assuming coverage of Aruba Networks (downgrade to Sell), Netgear (Neutral), Powerwave (Neutral), and Starent (downgrade to Neutral). We have also revised many of our estimates and price targets.

Industry view

In the long term, our coverage universe offers multiple opportunities for outsized growth, given a multiyear secular trend of accelerating bandwidth demand that is driving multiple product cycles around IP network convergence, 3G handsets, and wireless data. However, given softening near-term fundamentals and the group's 20.6X median P/E multiple, a 27% premium to the S&P 500, we are adopting a Neutral coverage view.

Risks

The primary downside risk to our view is that Street estimates have further downside, given many of our companies' significant enterprise and consumer exposure. The primary upside risk is that company-specific product cycles and 50% exposure to service provider spending, which we expect to be more resilient, drive EPS or multiple expansion.

Buy-rated stocks: CSCO, QCOM, RIMM

Our top idea is Cisco, as we believe the stock's 15.9X multiple is attractive given our view that estimates are near a bottom and that Cisco can achieve long-term growth of around 15%, as it benefits from (1) secular tailwinds, (2) solid execution, and (3) acquisition core competency. We also like QUALCOMM based on strong 3G handset growth, chipset share gains, and a likely legal resolution, and Research in Motion based on its solid execution and leadership in the fast-growing wireless email market.

Sell-rated stock: ARUN

We rate Aruba Networks Sell, as it is almost 100%-exposed to a weakening enterprise environment and will likely face more competition from Cisco.

TRANSITIONING COVERAGE

Acme Packet (APKT, Neutral)
Aruba Networks (ARUN, Sell)
Cisco Systems (CSCO, Conviction List Buy)
Juniper Networks (JNPR, Neutral)
Motorola (MOT, Not Rated)
Netgear (NTGR, Neutral)
Powerwave (PWAV, Neutral)
QUALCOMM (QCOM, Buy)
Research In Motion (RIMM, Buy)
Starent Networks (STAR, Neutral)

RATING CHANGES

CSCO – add to Conviction List – Buy-rated
JNPR – downgrade to Neutral from Buy
STAR – downgrade to Neutral from Buy
ARUN – downgrade to Sell from Neutral

UPCOMING EVENTS

11 a.m. ET, 03/27/08: Communication Technology Transfer of Coverage conference call hosted by Simona Jankowski and Thomas Lee; Dial-in: 877-208-2954 (Domestic), 973-528-0066 (International); Conference Entry Code: 289901

11 a.m. ET, 04/03/08: Conference call with Robert Lloyd, Senior Vice President, US, Canada and Japan, Cisco, hosted by Simona Jankowski; Dial-in: 877-208-2954 (Domestic), 973-528-0066 (International); Conference Entry Code: 217702

RELATED RESEARCH

03/26/08: Cisco Systems: Assuming coverage of Buy-rated CSCO, adding to Conviction List

03/26/08: Juniper Networks: Assuming coverage with Neutral: solid growth, premium valuation

Simona Jankowski, CFA
(415) 249-7437 | simona.jankowski@gs.com Goldman, Sachs & Co.

Thomas D. Lee
(212) 902-2066 | thomas.d.lee@gs.com Goldman, Sachs & Co.

Varun Marwah
(212) 357-5703 | varun.marwah@gs.com Goldman, Sachs & Co.

The Goldman Sachs Group, Inc. does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Customers in the US can receive independent, third-party research on companies covered in this report, at no cost to them, where such research is available. Customers can access this independent research at www.independentresearch.gs.com or call 1-866-727-7000. For Reg AC certification, see the text preceding the disclosures. For other important disclosures go to www.gs.com/research/hedge.html. Analysts employed by non-US affiliates are not required to take the NASD/NYSE analyst exam.

March 26, 2008

Americas: Communications Technology

Table of contents

Overview: Lowering industry coverage view to Neutral from Attractive, but still finding solid opportunities for absolute performance	2
Near-term industry outlook: Softening fundamentals	3
Valuation: Risk/reward more favorable post recent pullback	6
Long-term industry outlook: Tepid overall, but with pockets of strong growth	10
Service Provider segment capped by carrier capex growth in the mid-single digits	12
Enterprise network spending is converging to overall IT growth	13
Consumer segment growth (mostly handsets) is slowing	14
Acme Packet (APKT, Neutral, \$8 price target)	16
Aruba Networks (ARUN, Sell, \$4.50 price target)	18
Cisco Systems (CSCO, Conviction List Buy, \$29 price target)	20
Juniper Networks (JNPR, Neutral, \$25 price target)	23
Motorola, Inc. (MOT, Not Rated)	25
Netgear (NTGR, Neutral, \$21 price target)	26
Powerwave (PWAV, Neutral, \$2.30 price target)	28
QUALCOMM (QCOM, Buy, \$44 price target)	30
Research in Motion (RIMM, Buy, \$120 price target)	33
Starent Networks (STAR, Neutral, \$14 price target)	35
Disclosures	41

The prices in the body of this report are based on the market close of March 24, 2008.

Overview: Lowering industry coverage view to Neutral from Attractive, but still finding solid opportunities for absolute performance

We are lowering our Communications Technology industry coverage view to Neutral from Attractive, due to softening near-term fundamentals balanced by more attractive valuations. Our top stock ideas—**Cisco**, **Research in Motion**, and **QUALCOMM**—are characterized by secular tailwinds, solid execution, and reasonable valuations relative to their outsized growth prospects.

Softening near-term outlook... Near term, we are concerned about weakening trends in the Goldman Sachs CIO survey on enterprise network spending, as well as declining US communications equipment orders growth, a potential pause in network upgrades as capacity additions have finally caught up with bandwidth demand, and potential downside to margins as companies compete more aggressively for scarcer revenue opportunities.

... balanced by lower valuations. We believe our fundamental concerns are already largely reflected in the sector's valuations, as the GSTI Networking Index is down 21% year to date, and the average stock in our group is down 51% from its 52-week peak. Although we do not think we have reached the trough yet for the group, we do believe the risk/reward profile in our Buy-rated stocks offers opportunities for solid relative and absolute performance.

March 26, 2008

Americas: Communications Technology

Tepid long-term outlook but with pockets of strong growth. In the long term, while industry growth is tepid in the 5% area, the slice of the industry we cover offers double-digit growth opportunities due to its exposure to favorable secular trends such as the transition from circuit-switched to IP-based networks, 3G handsets, and wireless data.


Near-term industry outlook: Softening fundamentals

We are shifting our industry coverage view with a six-month horizon to Neutral from Attractive, as in the near term we see several signs that industry fundamentals are softening, which is not surprising given the macroeconomic slowdown. That being said, with the GSTI Networking Index down 21% year to date, we believe many of these risks are already priced into the stocks, as we discuss in more detail in the Valuation section below. In addition, there are multiple pockets of the industry that continue to show robust growth, as shown in Exhibit 13 below and discussed in more detail in our Cisco and Juniper transfer of coverage reports, published separately today. Therefore, we do not consider a Cautious industry coverage view appropriate at this point.

Some of the signs pointing to softening near-term industry fundamentals include:

1. **Weakening enterprise spending:** Our IT Spending Survey, which polls 100 Chief Information Officers (CIOs) on a bi-monthly basis, started noticeably weakening in October 2007 (Exhibit 1).

Exhibit 1: IT Spending Survey shows softening in CIO outlook for spending on networking
Responses by 100 CIOs to the question "In terms of your overall network equipment spending, what do you anticipate as the most likely scenario over the next 12 months?"




* Prior to October 2005, the question asked in the first half and second half of the year, respectively, was: "What do you anticipate the level of 2H [1H] network equipment spending at your company to be relative to the 1H [2H] spending"

Source: Goldman Sachs IT Spending Survey.


2. **US communications equipment orders rolling over:** US communications equipment orders, as reported by the US Department of Commerce durable goods report, have historically shown a high correlation with the GSTI Networking Index performance (Exhibit 2). Of note, orders dipped into negative territory for year-on-year growth in October 2007, and have remained flat to down since then (Exhibit 3).

March 26, 2008

Americas: Communications Technology

Exhibit 2: The GSTI Networking Index is 0.85 correlated with rolling average Communication Equipment orders

Source: US Census Bureau, FactSet, Goldman Sachs Research.

Exhibit 3: Declining Communications Equipment orders in late 2007/early 2008 reflect a softening environment

Source: US Census Bureau, Goldman Sachs Research.


3. **Less urgency to upgrade core networks in the near term:** Although bandwidth demand growth remains extremely robust, capacity additions to the core of the network actually exceeded the rate of demand growth in a meaningful way for the first time in eight years (Exhibit 4). Even though accelerating bandwidth demand growth will likely continue to drive robust growth in certain segments tied to next-generation network buildouts in the long term, we think in the near term carriers have some breathing room to potentially delay purchases if the environment were to deteriorate, given their brisk pace of capacity additions in the last three years.

March 26, 2008

Americas: Communications Technology

Exhibit 4: Core capacity additions have finally caught up with bandwidth demand growth

Core capacity growth (as measured by core router port shipments) and bandwidth demand growth (as measured by broadband subscriber growth and bandwidth use per subscriber)



Source: Dell'Oro, TeleGeography, Goldman Sachs Research estimates.


- 4. Downside risk to margins:** Communications Technology company margins remain very close to their all-time peak levels of 2005, and far above prior downturns in 1993 and 1998, even setting aside the tech and telecom bust of 2001-2002 (Exhibit 5). We believe this suggests that there is risk to Street estimates not only from slower-than-expected top-line growth, as suggested in our first three points above, but more importantly from deteriorating margins, as companies start to price more aggressively and spend more on sales and marketing in an attempt to drive the top line.

March 26, 2008

Americas: Communications Technology

Exhibit 5: CommTech margins have downside risk

net profit margins for the S&P 500 Technology and Communications Equipment constituents



Source: FactSet, Goldman Sachs Research.

To put the above data points on IT spending, US communications orders, and core network capacity additions in context, we show below the exposure for the companies in our coverage to the key geographic and vertical end markets (Exhibit 6).

Exhibit 6: Significant Enterprise and Service Provider exposure

sales split by key vertical

Company	Geographic split					Vertical split			
	North America	Europe	APAC	Japan	ROW	Enterprise	Service Provider	SMB	Consumer
Acme Packet	46%				54%	5%	95%		
Aruba Networks	67%	16%	10%		7%	100%	0%		0%
Cisco	55%	21%	10%	4%	10%	40%	35%	23%	3%
Juniper Networks	47%	32%	21%			36%	64%		
Motorola	51%	13%	16%		20%	21%	16%		63%
Netgear	38%	52%	10%		0%	0%	22%	36%	42%
Powerwave	28%	45%	26%		0%		100%		
Qualcomm	13%		52%	18%	17%	na	na	na	na
RIMM	84%				16%	73%			27%
Starent	76%		9%	13%	2%		100%		
Median	49%	27%	13%	13%	10%	36%	50%	30%	27%

Source: Company data, Goldman Sachs Research estimates.

Valuation: Risk/reward more favorable post recent pullback

Our view that industry fundamentals are softening in the near term is already largely reflected in the 21% year-to-date decline in the GSTI Networking Index—the largest decline among all Technology sectors. Although this is likely not yet the bottom, this pullback has created a more favorable risk/reward profile for the group, balancing our near-term caution on the fundamentals and supporting our Neutral sector coverage view. Using our P/E-based valuation methodology,

March 26, 2008

Americas: Communications Technology

we see several opportunities for both relative and absolute performance over the next six months, including Cisco, Research in Motion, and QUALCOMM.

P/E-based valuation methodology

We value growth stocks in our coverage based on a growth-adjusted target P/E multiple, applied to our CY2009 EPS estimate (Exhibit 7). In particular, we first derive an estimated five-year EPS growth rate based on the company's end-market growth, new market penetration, market-share trajectory, and margin trends. We then apply a target price-to-earnings-to-growth (PEG) multiple to that growth rate, ranging from 0.9X to 1.2X based on our view of the quality, predictability, and volatility of earnings. This allows us to obtain a target P/E multiple.

We value turnaround stocks in our coverage based on a target enterprise value to sales (EV/S) multiple applied to our CY2009 sales estimate. Our target EV/S multiple is based on our view of how the company's progress in its turnaround and likely eventual outcome positions it relative to its historical multiple range and its peers' valuations.

Based on our analysis, we see very little upside in our coverage universe over the next six months, which we believe is not high enough to justify an Attractive sector view given the well above-average risk and volatility in the group. Our top Buy-rated names include Cisco, Research in Motion, and QUALCOMM, and our Sell-rated name is Aruba Networks (Exhibit 7).

Exhibit 7: Price targets and methodology for our covered companies

we use PEG for growth companies and EV/S for turnaround companies

	GS Stock	Price	Time	Price	Market	CY08	CY08	CY08	5-yr EPS	Target	CY08	Upside/	
	Ticker	Rating	Target	Frame	3/24/08	Cap	EPS	P/E	PEG	CAGR	PEG	Target P/E	Downside
Growth (PEG)													
Acme Packet Inc.	APKT	Neutral	\$8	6 months	\$7.80	\$513	\$0.43	18.0 x	0.9 x	20%	0.9 x	18.0 x	0%
Aruba Networks Inc.	ARUN	Sell	\$4.5	6 months	\$5.64	\$527	\$0.14	41.7 x	1.2 x	35%	0.9 x	32.9 x	-21%
Cisco Systems Inc.	CSCO	CL Buy	\$29	6 months	\$25.64	\$160,583	\$1.61	15.9 x	1.1 x	15%	1.2 x	18.0 x	13%
Juniper Networks Inc.	JNPR	Neutral	\$25	6 months	\$25.50	\$14,316	\$1.13	22.6 x	1.1 x	20%	1.1 x	22.0 x	-3%
NETGEAR Inc.	NTGR	Neutral	\$21	6 months	\$20.54	\$700	\$1.72	11.9 x	1.0 x	12%	1.0 x	12.0 x	1%
QUALCOMM Inc.	QCOM	Buy	\$44	6 months	\$39.89	\$68,651	\$2.15	18.6 x	1.0 x	18%	1.1 x	20.7 x	12%
Research In Motion Ltd.	RIMM	Buy	\$120	6 months	\$111.83	\$64,249	\$4.16	28.9 x	1.2 x	23%	1.2 x	28.7 x	7%
Starent Networks Corp.	STAR	Neutral	\$14	6 months	\$14.64	\$843	\$0.48	30.6 x	1.1 x	28%	1.0 x	29.3 x	-4%
							CY08 Sales	CY08 EV/S	CY08 Target EV/S				
Turnaround (EV/S)													
Motorola Inc.	MOT	Not Rated			\$9.69	\$23,872	\$37.875	0.50 x					
Powerwave Technologies Inc.	PWAV	Neutral	\$2.3	6 months	\$2.53	\$344	\$949.49	0.68 x				0.65 x	-9%

For important disclosures, please go to <http://www.gs.com/research/hedge.html>.

Our growth stock valuation methodology is described above. We value turnaround stocks based on a target CY2009 enterprise value to sales (EV/S) multiple, based on our view of how the company's progress in its turnaround and likely eventual outcome positions it relative to its historical multiple range and its peers' valuations.

Source: Company data, FactSet, Goldman Research estimates.

March 26, 2008

Americas: Communications Technology


Exhibit 8: Communications Equipment risks

Company	Risks
Acme Packet Inc.	(1) A tougher competitive environment as larger competitors such as Cisco, Juniper and Sonus have broader product lines, deeper pockets, and the capability to integrate SBC functionality into their other offerings; (2) high customer concentration; and (3) lumpy quarter-to-quarter patterns given the early stage of the technology deployment.
Aruba Networks Inc.	(1) A less severe IT spending environment than we expect; (2) lack of increased competitive pressure from Cisco; and (3) successful implementation of Aruba Networks' 2-tier distribution system, resulting in greater than expected share gains and operating margin expansion.
Cisco Systems Inc.	(1) A slowdown in IT spending; (2) a slowdown in carrier spending; (3) deceleration in emerging markets; (4) declines in Cisco's router and switch market share; (5) potential transition from pro forma to GAAP estimates by the Street.
Juniper Networks Inc.	(1) A slowdown in carrier spending; (2) weakening in Juniper's competitive positioning; (3) execution in ramp of new products; (4) failure to achieve operating leverage; (5) high bar for the stock given premium multiple and high investor expectations; (6) transition from pro forma to GAAP estimates by the Street.
Motorola Inc.	Not Rated (NR)
NETGEAR Inc.	(1) Delays in distributing products to key markets; (2) increasing competition from established players with deeper pockets; and (3) lack of operating leverage.
Powerwave Technologies Inc.	(1) Continued deterioration in wireless infrastructure market; (2) weak spending patterns at large customers; and (3) pricing pressure
QUALCOMM Inc.	(1) Changes to the royalty rate, (2) increased WCDMA chipset competition, (3) CDMA/WCDMA handset market weakness, (4) a further deterioration in the macro environment.
Research In Motion Ltd.	(1) Lower enterprise IT spending, (2) Increasing competition, (3) Increased investment may reduce operating leverage, (4) network operating center (NOC) scalability.
Starent Networks Corp.	(1) Increased competition; (2) high customer concentration; (3) lumpy order patterns.

Source: Goldman Sachs Research estimates.

Stocks are not yet trading at trough multiples

Although valuations overall have become much more attractive, with the median stock in our coverage down 51% from its 52-week highs, we are not ready to call this the bottom, for two reasons. First, we see further risk to estimates based on the revenue and margin concerns discussed above. Second, the stocks are not yet trading at trough multiples (Exhibit 9) and in several cases are still well off their 52-week lows (Exhibit 10).

Exhibit 9: Closer to trough multiples, but not there yet
P/E multiples relative to the S&P 500 for key tech sectors

Source: FactSet, Goldman Sachs Research.

Exhibit 10: Several stocks are still well off their lows
stocks relative to their 52-week highs and lows

Ticker	Price	52-Week High		52-Week Low	
		Price	% Chg	Price	% Chg
APKT	\$7.80	\$16.59	(53%)	\$6.85	14%
ARUN	\$5.64	\$23.85	(76%)	\$4.63	22%
CSCO	\$25.64	\$34.24	(25%)	\$21.77	18%
JNPR	\$25.50	\$37.95	(33%)	\$18.08	41%
MOT	\$9.69	\$19.68	(51%)	\$8.98	8%
NTGR	\$20.54	\$41.33	(50%)	\$18.58	11%
PWAV	\$2.53	\$7.64	(67%)	\$2.12	19%
QCOM	\$39.89	\$47.72	(16%)	\$35.17	13%
RIMM	\$111.83	\$137.01	(18%)	\$42.93	160%
STAR	\$14.64	\$31.67	(54%)	\$10.00	46%
Median			(51%)		19%


Source: FactSet, Goldman Sachs Research.

March 26, 2008

Americas: Communications Technology

In addition, investor sentiment appears relatively neutral, with short interest at relatively average levels, suggesting that hedge funds are not positioned for another round of misses in the first-quarter reporting season (Exhibit 11).

Exhibit 11: Short interest suggests investors are not positioned for another round of cuts
Nasdaq short interest levels in our coverage universe



Source: Nasdaq, Goldman Sachs Research.

March 26, 2008

Americas: Communications Technology

Exhibit 12: Communications Equipment valuation table

Ticker	GS Stock Rating	Shares O/S (M)	Price 3/24/2008	Market Cap (\$M)	EPS (ex ESO)		EPS (incl ESO)		P/E (ex ESO)		P/E (incl ESO)		
					CY08E	CY09E	CY08E	CY09E	CY08E	CY09E	CY08E	CY09E	
Growth													
Micro-Cap													
Acme Packet Inc.	APKT	Neutral	65.8	\$7.80	\$513	\$0.40	\$0.49	\$0.34	\$0.43	19.5 x	15.8 x	23.0 x	18.0 x
Aruba Networks Inc.	ARUN	Sell	93.5	\$5.64	\$527	\$0.06	\$0.25	-\$0.11	\$0.14	90.2 x	22.6 x	NM	41.3 x
Small-Cap													
NETGEAR Inc.	NTGR	Neutral	34.1	\$20.54	\$700	\$1.78	\$1.99	\$1.52	\$1.72	11.5 x	10.3 x	13.5 x	11.9 x
Starent Networks Corp.	STAR	Neutral	57.6	\$14.64	\$843	\$0.62	\$0.70	\$0.40	\$0.48	23.8 x	20.8 x	36.5 x	30.6 x
D-Link Corp.	625487	Neutral	547.1	\$50.80	\$27,792	\$0.15	\$0.21	\$0.15	\$0.19	NM	NM	NM	NM
Large-Cap													
Cisco Systems Inc.	CSCO	CL Buy	6263.0	\$25.64	\$160,583	\$1.59	\$1.85	\$1.37	\$1.61	16.1 x	13.9 x	18.7 x	15.9 x
Juniper Networks Inc.	JNPR	Neutral	561.4	\$25.50	\$14,316	\$1.07	\$1.34	\$0.88	\$1.15	23.9 x	19.1 x	29.1 x	22.3 x
QUALCOMM Inc.	QCOM	Buy	1721.0	\$39.89	\$68,651	\$2.09	\$2.33	\$1.89	\$2.15	19.1 x	17.1 x	21.1 x	18.6 x
Research In Motion Ltd.	RIMM	Buy	574.5	\$111.83	\$64,249	\$3.09	\$4.16	\$3.09	\$4.16	36.2 x	26.9 x	36.2 x	26.9 x
Group Median										21.6 x	18.1 x	23.0 x	20.4 x
Micro-cap Growth Median										54.9 x	19.2 x	23.0 x	29.7 x
Small-cap Growth Median										17.7 x	15.5 x	25.0 x	21.2 x
Large-cap Growth Median										21.5 x	18.1 x	25.1 x	20.4 x
Mature													
L.M. Ericsson Telephone Co. (ADS)	ERIC	Sell	1589.6	\$18.48	\$29,376	\$1.31	\$1.65	\$1.31	\$1.65	14.1 x	11.2 x	14.1 x	11.2 x
Nokia Corp. (ADR)	NOK	Neutral	3885.7	\$30.30	\$117,736	\$2.56	\$2.79	\$2.56	\$2.79	11.8 x	10.9 x	11.8 x	10.9 x
Group Median										13.0 x	11.0 x	13.0 x	11.0 x
Turnaround													
Motorola Inc.	MOT	Not Rated	2463.6	\$9.69	\$23,872	\$0.24	\$0.63	\$0.11	\$0.61	40.4 x	15.4 x	88.1 x	15.9 x
Powerwave Technologies Inc.	PWAV	Neutral	135.8	\$2.53	\$344	\$0.05	\$0.27	\$0.02	\$0.24	51.7 x	9.4 x	136.7 x	10.6 x
Alcatel-Lucent (ADS)	ALU	Sell	2257.1	\$5.49	\$12,391	\$0.25	\$0.43	\$0.25	\$0.43	22.0 x	12.9 x	22.0 x	12.9 x
Group Median										40.4 x	12.9 x	88.1 x	12.9 x
Super Large Cap Tech													
Cisco Systems Inc.	CSCO	CL Buy	6,263.0	\$25.64	\$160,583	\$1.59	\$1.85	\$1.37	\$1.61	16.1 x	13.9 x	18.7 x	15.9 x
Dell Inc.	DELL	Buy	2266.0	\$20.56	\$46,589	\$1.48	\$1.55	\$1.48	\$1.73	13.9 x	13.3 x	13.9 x	11.9 x
Intel Corp.	INTC	Buy	5988.0	\$22.13	\$132,514	\$1.39	\$1.79	\$1.24	\$1.65	15.9 x	12.4 x	17.9 x	13.4 x
International Business Machines Corp.	IBM	Buy	1412.9	\$119.06	\$168,220	\$8.32	\$9.35	\$8.10	\$9.15	14.3 x	12.7 x	14.7 x	13.0 x
Nokia Corp. (ADR)	NOK	Neutral	3,885.7	\$30.30	\$117,736	\$2.56	\$2.79	\$2.56	\$2.79	11.8 x	10.9 x	11.8 x	10.9 x
Microsoft Corp.	MSFT	Not Rated	9503.0	\$29.17	\$277,203	\$1.98	\$2.16	\$1.96	\$2.20	14.8 x	13.5 x	14.9 x	13.3 x
Group Median										14.5 x	13.0 x	14.8 x	13.1 x
Universe Median										16.1 x	13.5 x	18.7 x	13.4 x
Major Indices													
S&P 500	SPX			\$1,349.88				\$78.00	\$83.00			17.3 x	16.3 x

For important disclosures, please go to <http://www.gs.com/research/hedge.html>.

Primary coverage responsibilities: APKT, CSCO, JNPR by Simona Jankowski; ARUN, MOT, NTGR, PWAV, QCOM, RIMM, STAR by Thomas Lee; DELL, IBM by David Bailey; Intel by James Covello; Microsoft by Sarah Friar; Alcatel-Lucent, L.M. Ericsson and Nokia by Tim Boddy; D-Link by Joey Cheng. Estimates for D-Link have been translated to US\$ from NT\$. MOT estimates are consensus estimates.

Source: Company data, Goldman Sachs Research estimate, FactSet, First Call.

Long-term industry outlook: Trepid overall, but with pockets of strong growth

Trepid growth overall... In our view, the Communications Technology industry, spanning all segments including telecom equipment, data networking, and handsets, is likely to grow at an uninspiring mid-single-digit growth rate in the long term, or relatively in line with GDP. As described in more detail below, our assumptions include 6% growth in the Consumer segment of the industry (mostly handsets), which drives around 55% of total

March 26, 2008

Americas: Communications Technology

sales, 5% growth in the Service Provider segment, which drives another 30%, and 7% growth in the Enterprise segment, which drives the remaining 14%.

... with pockets of faster growth driven by secular tailwinds. Within the tepid overall growth rate, there are several pockets of robust double-digit growth (Exhibit 13); at the highest level, these include:

- **Transition from circuit-switched voice to next-generation IP-based networks:** Carriers are transitioning away from their legacy circuit-switched voice networks to next-generation IP-based networks, as that enables them to lower their operating costs and drive higher revenues by offering IP-based services and not just basic connectivity. This transition is driving robust 20%-plus growth in service provider routers and switches, as well as in new types of equipment such as voice over IP (VoIP) and IPTV that provide richer, cheaper services by leveraging the IP network. This trend favors Cisco and Juniper, the leaders in service provider routers and other IP equipment.
- **Transition from 2G/2.5G to 3G handsets:** As the handset market transition from second- (2G) to third-generation (3G) technologies, we are forecasting the WCDMA (3G) segment to grow at a five-year unit CAGR of 27% from 2007 to 2012, well above the overall handset unit growth of 9%. We expect the mix of WCDMA handsets, currently around 16% of total handset shipments, to increase to 34% in 2012. By contrast, we expect the GSM (2G) market, which currently represents 65% of the total handset market, to decline to 50%. This trend favors QUALCOMM, as we estimate that WCDMA represents 48% of the company's revenues.
- **Growth in wireless data:** We are at the early stages of seeing significant growth in wireless data, driven by the intense focus from wireless operators to increase their data subscribers through higher subsidies for smartphones, the availability of more affordable smartphones, and increased consumer appetite for mobile applications such as Internet access, wireless e-mail, GPS applications, and gaming. The growth in wireless data is driving significant growth in the smartphone market, which we forecast to grow at a five-year unit CAGR of 26% to reach 362 million units in 2012 (20% of the handset market), up from 114 million units in 2007 (10% of the handset market). Based on our price erosion assumptions, this translates to a five-year revenue CAGR of 13%, reaching \$66 billion in 2012 or 36% of the total handset market (vs. 23% in 2007). This trend benefits QUALCOMM due to its dominant position in core 3G technology, which is at the heart of most smartphones, Research in Motion due to its leading share in smartphones, and Starent due to its solid positioning in wireless data infrastructure equipment.

March 26, 2008

Americas: Communications Technology

Exhibit 13: We expect mid-single-digit growth for the Communications Technology industry in the long term
market size, expected growth, competition, and drivers

Market Segment	% of Total	% of Segment	CY06 Market Size (\$mn)	06-10E CAGR	Competitors	Driver
Service Provider	31%	% of Service Provider	\$79,927	6%		
Wireless Radio Access and Backhaul Equipment	18%	57%	\$45,856	1%	ERIC, Nokia-Siemens, ALU, Huawei	Declining spending in GSM and CDMA offset by WCDMA investments
Optical Network Hardware	5%	15%	\$11,954	0%	SONET/SDH: ALU, Huawei, TLAB; WDM: ALU, CSCO, Huawei, NT	Transition from SONET/SDH (2/3 of market) to WDM (1/3 of market)
Service Provider Routers and Switches	4%	12%	\$9,820	20%	CSCO, JNPR, ALU, Huawei	Major IP network transformation projects initiated by carriers
Broadband Aggregation Equipment	3%	10%	\$7,987	6%	DSLAM: ALU, Huawei, ZTE, MSAP; NSN, ZTE	Replacement of ATM DSLAM by IP DSLAM and deployment in Asia
Service Provider VoIP and IMS Equipment	1%	4%	\$3,304	20%	VoIP: CSCO, SONS, NT; Softswitch: NT, Nokia-Siemens, Huawei	Migration of circuit-switched voice networks to next-gen Packet networks
IPTV Equipment	0%	1%	\$1,005	29%	MOT, CSCO, ALU, Thomson	Extensive rollout by Service providers to capture triple play revenue
Enterprise	14%	% of Enterprise	\$36,982	7%		
Enterprise Routers and Switches	8%	57%	\$21,054	6%	Router: CSCO, JNPR, Huawei, Switch: CSCO, HPQ, NT, 3Com	Enterprise demand for secure and better connectivity
Enterprise Telephony Equipment	3%	24%	\$8,908	7%	Avaya, CSCO, NT	Migration from circuit switches to IP PBX packet switches
Network Security Appliances and Software	2%	14%	\$5,211	9%	CSCO, JNPR, Check Point	Increasing spending on protection against loss of data and security threats
Wireless LAN Equipment	1%	5%	\$1,810	18%	CSCO (including Linksys), D-Link, NTGR, MOT	Rollout of WiFi hotspots and adoption of wireless LANs by Enterprises
Consumer	55%	% of Consumer	\$146,076	6%		
Handsets	53%	96%	\$138,938	6%	NOK, Samsung, MOT, SNE, LG	Replacement demand in mature markets, new subscribers in emerging markets
- Smartphones			\$25,340	23%	NOK, RIM, Apple, Palm, MOT	Increasing emphasis by operators on driving wireless data growth; Shift in consumer interest on devices to applications/usability vs industrial design.
Broadband CPE	2%	4%	\$6,137	7%	Cable: Arris, MOT, CSCO DSL: Thomson, ZyXEL	Upgrade in broadband access networks for more bandwidth and services
Total	100%		\$261,984	6%		

Source: Infonetics, Goldman Sachs Research estimates.


Service Provider segment capped by carrier capex growth in the mid-single digits

Fundamentally, the Service Provider segment of the industry, which accounts for around a third of industry revenues, is driven by telecommunications carrier capital expenditures. With the carrier capex/sales ratio stable at around 17%, which is in line with the long-term trend and go-forward targets, we expect the capex-driven portion of the industry to grow in line with the telecommunications industry, or in the mid-single-digit range (Exhibit 14).

March 26, 2008

Americas: Communications Technology

Exhibit 14: Carrier capex growth is likely to remain in the low- to mid-single-digit range
 carrier capex/sales ratio and year-on-year capex growth



Source: Company data, Goldman Sachs Research estimates.


Enterprise network spending is converging to overall IT growth

The Enterprise segment of the Communications Technology industry, which accounts for around 14% of total revenues, is driven largely by IT spending, as large enterprises and smaller businesses spend on building and upgrading their networks (Exhibit 15). Although in the last 18 years, the correlation between overall IT spending growth and network spending growth was relatively low at 73%, in the last five years it has increased to 98%. Therefore, we expect growth in the Enterprise segment of the Communications Technology industry to converge to that of overall IT spending, or the mid- to high-single-digits over the next few years.

March 26, 2008

Americas: Communications Technology

Exhibit 15: Network spending growth is converging with IT spending growth at ~7%
 year-on-year change in overall IT spending and network equipment spending



Source: IDC.


Consumer segment growth (mostly handsets) is slowing

The Consumer segment of the Communications Technology industry, which accounts for around 55% of total revenues, is driven largely by handsets, and to a lesser degree by customer premise equipment (CPE) such as set-top boxes. We believe handset unit growth has substantially slowed from the 15%-30% range in the last five years to the 10% level in the next few years. Given average selling price declines of around 5% annually, we expect revenue growth for the segment to be around 5% going forward (Exhibit 16).

March 26, 2008

Americas: Communications Technology

Exhibit 16: Handset revenue growth slowing to the mid-single-digits
handset unit forecast, handset unit growth and handset revenue growth



Source: Goldman Sachs Research estimates.


March 26, 2008

Americas: Communications Technology

Acme Packet (APKT, Neutral, \$8 price target)

Simona Jankowski is assuming primary coverage of Acme Packet from Brantley Thompson, who is moving to the buy-side, and maintaining a Neutral rating. Our six-month price target of \$8 (down from \$10 previously) is based on 18.0X our CY2009 EPS (including ESO) estimate of \$0.43. Our multiple is based on a target PEG multiple of 0.9X and a long-term EPS growth rate of 20%. Our EPS growth expectations are based on (1) the rapid growth of the session border control (SBC) market, where Acme Packet has the leading share, (2) the company's ability to defend its share based on its technological differentiation, and (3) stable operating margins despite increasing competitive pressures, as Acme Packet's robust top-line growth helps drive operating leverage.

Exhibit 17: Acme Packet at a glance



Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- Rapidly growing addressable market.** Acme Packet has approximately 50% share of the \$200 million session border control (SBC) market. SBC equipment is deployed by carriers to control and manage sessions at the border points between networks (a session is a real-time communication on the network such as video on demand or voice over IP). Demand for this type of equipment is growing rapidly, due to (1) an increasing number of sessions on carrier networks, as carriers deploy more advanced new services such as voice over IP and IPTV, (2) a greater number of border points where one IP network meets another, and (3) security and law enforcement requirements. Therefore, we expect the session border control (SBC) market to grow at a robust CAGR in the 30% area over the next 3-5 years.
- Technological differentiation helps defend share.** Acme Packet's products offer a strong feature set, broad interoperability with major vendors, and the reliability and scalability required by service providers and cable operators. The company's technological differentiation has earned it a marquee list of more than 500 customers, including 24 of the top 25 (and 82 of the top 100) global service providers, and 7 of the top 10 cable operators. Key customers include EarthLink, Global Crossing, Telefonica, Vodafone, British Telecom, France Telecom, China Unicom, NTT Telecom, KDDI, and Chungwa Telecom. Acme Packet products are distributed by every major telecom equipment vendor, including Ericsson, Sonus Networks, Alcatel, Lucent, Nortel, and Motorola, as many vendors do not offer their own SBC solutions. Over time, we see an

March 26, 2008

Americas: Communications Technology

increasingly competitive environment, as larger vendors such as Cisco and Juniper have embedded SBC functionality into their routers, and Sonus Networks has embedded it into their softswitch products. However, given that routers are not found in many of the locations in which SBC is needed, and the significant complexity and hardware intensity of SBC equipment, we expect there to be a market for stand-alone SBC equipment for at least a few years.

- **Stable operating margins.** We estimate that GAAP operating margins will be in the 21%-22% range during 2008-2010, a healthy level for a company of Acme Packet's size. Despite the likely increasingly competitive environment, we expect the company will maintain its margins as its robust top-line growth helps drive operating leverage.

Valuation

We rate Acme Packet Neutral with a six-month price target of \$8. We arrive at our six-month price target by applying an 18.0X multiple to our fiscal 2009 EPS (including ESO) estimate of \$0.43. Our multiple is based on a target PEG multiple of 0.9X and a long-term EPS growth rate of 20%.

Acme packet's current CY2009E multiple of 18.0X represents a 12% discount to the median growth company P/E in our group of 20.6X, and a 11% premium to S&P 500 P/E of 16.3X.

Catalysts

(1) Acme Packet's quarterly results. (2) Continued product traction at major carriers. Revenue driver for the company is shifting from new customer additions to deeper penetration of the existing 500-plus installed base.

Risks

Risks include (1) a tougher competitive environment as larger competitors such as Cisco, Juniper, and Sonus have broader product lines, deeper pockets, and the capability to integrate SBC functionality into their other offerings; (2) high customer concentration; and (3) lumpy quarter-to-quarter patterns given the early stage of the technology deployment.

Estimate changes

We are transitioning to GAAP EPS as our primary valuation metric. GAAP estimates for Acme Packet differ from pro forma estimates primarily in that they exclude share-based compensation, net of tax. We believe that consensus estimates for Acme Packet are currently pro forma, so we have shown both our GAAP and pro forma estimates in Exhibit 18 for comparison purposes.

March 26, 2008

Americas: Communications Technology

Exhibit 18: Acme Packet


(\$ in millions)	2Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$35	\$145	\$180	\$218
GS - current	\$35	\$145	\$180	\$218
Consensus	\$34	\$144	\$183	\$230
EPS				
GS (GAAP) - previous	\$0.08	\$0.33	\$0.43	\$0.51
GS (GAAP) - current	\$0.09	\$0.34	\$0.43	\$0.52
GS (ex-ESO) - previous	\$0.10	\$0.40	\$0.49	\$0.59
GS (ex-ESO) - current	\$0.10	\$0.40	\$0.49	\$0.59
Consensus (ex-ESOs)	\$0.09	\$0.40	\$0.52	-

Source: FactSet, Goldman Sachs Research estimates.

Aruba Networks (ARUN, Sell, \$4.50 price target)

Thomas Lee is assuming primary coverage of Aruba Networks from Brantley Thompson, who is moving to the buy-side. We are downgrading the stock to Sell from Neutral and reducing our six-month price target to \$4.50 from \$5. Our six-month price target of \$4.50 is based on 32.9X our CY2009 GAAP EPS estimate of \$0.14. Our multiple is based on a target PEG multiple of 0.9X and a long-term EPS growth rate of 35%.

Although we believe Aruba Networks has a compelling offering and its products may be technically superior to the competition's in certain aspects, we believe Aruba faces considerable amount of risk during an economic downturn. Specifically, in addition to slowing end-market demand, we believe that customers could increasingly choose Cisco over Aruba, given Cisco's attractive bundled pricing and the perceived higher risk of going with a smaller vendor during uncertain economic times.

Exhibit 19: Aruba Networks at a glance

Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

March 26, 2008

Americas: Communications Technology

Investment view

- **Large enterprise exposure poses significant risk in current environment; we expect next several quarters to be challenging.** In our recent survey, CIOs appear to be increasingly bearish on the current macro outlook, and we expect this to have a tightening effect on IT budgets, at least in the first half of the year. This clearly weighed on Aruba's recent results, when the company missed its original revenue guidance for 2QFY08 (January 2008) by roughly 20% (negatively preannounced on February 7). The company cited weakness in certain verticals in both the US (federal and retail) and Europe (retail). Although the company believes the federal vertical is likely to return to normal levels in 3QFY08 as weakness was driven by a delay in spending, we believe that it is likely that enterprise weakness will continue for several quarters.
- **Longer-term concern: Macro slowdown could pave the way for share gains by Cisco.** On its earnings call, Cisco talked about being opportunistic in gaining share in certain markets during a potential economic slowdown. During the last economic downturn in 2001-2003, Cisco used its leading position to gain significant share in areas such as Ethernet switching. For instance, from 2000 to 2003, Cisco gained almost 800 bp of revenue share in Ethernet switching, largely at the expense of 3COM (lost 140 bp), Alcatel Lucent (lost 70 bp), Extreme (lost 85 bp), and Nortel (lost almost 400 bp). In the current slowdown, we believe Cisco could offer enterprises attractive bundling packages, essentially giving away its WLAN portion, which would make the competitive environment more challenging for Aruba. In addition, we believe that during uncertain economic times, IT managers are less likely to choose the solution of a small vendor over that of an established partner that they can have greater confidence will support the product for many years.
- **Product offering, while compelling, may not be enough to sway enterprises to choose Aruba over Cisco.** Although Aruba has an attractive solution (strong centralized management and security feature set) from a pure WLAN standpoint, Cisco's ability to provide a complete end-to-end solution (integrating security, wireless, VPN, VoIP, routing, and switching), gives it a significant advantage over its competitors. This is particularly the case for enterprises where WLAN is only a portion of the overall networking solution. On the positive front, we believe that Aruba is gaining share from Motorola's Symbol business, given its superior software and security technology.

Valuation

We are lowering our six-month price target of \$4.50 (from \$5), which is based on 32.9X our CY2009 GAAP EPS estimate of \$0.14. Our target P/E is based on a PEG multiple of 0.9X (low end of range in our coverage universe) and a long-term EPS growth rate of 35%. This robust growth expectation is driven by roughly 20% market growth, as well as significant operating margin expansion from our 5% estimate in CY2009 into the double-digits, as Aruba grows its revenue base and improves its channel leverage.

Aruba's current CY2009E multiple of 41.7X represents a 102% premium to the median growth company P/E in our group of 20.6X, and a 156% premium to S&P 500 P/E of 16.3X.

Risks

Risks include (1) a less severe IT spending environment than we expect; (2) lack of increased competitive pressure from Cisco; and (3) successful implementation of Aruba

March 26, 2008

Americas: Communications Technology

Networks' 2-tier distribution system, resulting in greater-than-expected share gains and operating margin expansion.

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for Aruba differ from pro-forma estimates primarily in that they exclude share-based compensation, net of tax. We believe that Consensus estimates for Aruba are currently pro forma, so we have shown changes to both our GAAP and pro forma estimates in Exhibit 20 for comparison purposes.

Exhibit 20: Aruba Networks

(\$ in millions)	3Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$43	\$176	\$227	\$288
GS - current	\$41	\$171	\$207	\$250
Consensus	\$43	\$177	\$236	\$288
EPS				
GS (GAAP) - previous	(\$0.07)	(\$0.12)	\$0.14	\$0.27
GS (GAAP) - current	(\$0.07)	(\$0.15)	\$0.05	\$0.21
GS (ex-ESO) - previous	(\$0.01)	\$0.06	\$0.30	\$0.38
GS (ex-ESO) - current	(\$0.02)	\$0.05	\$0.18	\$0.31
Consensus (ex-ESOs)	(\$0.01)	\$0.06	\$0.25	NA

Source: FactSet, Goldman Sachs Research estimates.

Cisco Systems (CSCO, Conviction List Buy, \$29 price target)


Please also refer to our detailed Cisco report published separately today: "Assuming coverage of Buy-rated CSCO, adding to Conviction List," March 26, 2008.

Simona Jankowski has assumed primary coverage of Cisco from Brantley Thompson, who is moving to the buy-side. We have added the Buy-rated stock to our Conviction List, and raised our six-month price target to \$29 from \$28. Our price target is based on 18.0X our CY2009 GAAP EPS estimate of \$1.61, with the target P/E based on a PEG multiple of 1.2X (high end of the range in our coverage universe given the higher quality and predictability of Cisco's earnings) and a long-term EPS growth rate of 15%. Our growth expectations are based on (1) secular tailwinds such as increased network intelligence, the transition to IP-based networks and accelerating bandwidth growth, (2) consistent execution in establishing and/or maintaining a leadership position in its key markets, and (3) core competency in creating demand and revenue synergies through acquisitions.

With the recent reset of the Street's estimates for Cisco, we believe the market will refocus on the solid earnings growth story, allowing the multiple to expand to more fairly reflect Cisco's above-average growth. Our confidence that consensus estimates for Cisco have bottomed is bolstered by our latest IT Spending Survey, which shows improvement in CIOs' outlook for spending on Cisco products.

March 26, 2008

Americas: Communications Technology

Exhibit 21: Cisco Systems at a glance

Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- Near-term macro concerns already in our estimates and likely in the stock.** Cisco shares are down 27% from their peak in November, compared with the S&P 500 down 13% and the Nasdaq down 20%. With the company guiding well below the Street for its April quarter, and reducing full-year revenue growth expectations toward the low end of its original 13%-16% guidance, we believe the macro concerns are now reflected in both our estimates and in the stock price. In particular, our estimates for the remainder of CY2008 assume further deceleration in the routing/switching markets (together 55% of sales) from 10% year-on-year growth in the January 2008 quarter to 3% in the seasonally slowest October 2008 quarter, when we expect the macro trends to bottom. We view these estimates as conservative enough, given the historical growth rate of the combined routing/switching market of 13% annually over the last three years, and our expected forward three-year growth rate for these markets of around 10%. In addition, our checks suggest continued strength in Cisco's small-and-medium business (SMB) segment (23% of sales) and in emerging markets (20% of sales).
- Improving CIO outlook on Cisco in our latest IT Spending Survey.** Our Technology Research team's IT Spending Survey, which polls 100 CIOs on a bi-monthly basis, showed a meaningful rebound in the outlook for spending on Cisco products in its February reading. In particular, 53% of respondents expect their spending on Cisco products to increase over the next 12 months—an increase from the 20% level in October and 34% level in December. The 53% positive response rate is back in the 50%-70% range that we consider healthy for Cisco's outlook—a range that was maintained in the two-year period from June 2005 through August 2007. In addition, the percentage of respondents who expect their Cisco spending to decline remains below 10%, another healthy indication. We view this as a favorable near-term indicator for Cisco's Enterprise business, as our survey results have historically been a reasonably good predictor of Cisco's fundamentals on the enterprise side, which drives 40% of sales.
- Secular tailwinds, solid execution, and core competency in acquisitions.** We expect Cisco to grow its top- and bottom-line in the midteens range over the next 3-5 years, as the company is well-positioned to benefit from secular tailwinds in all of its major businesses, including (1) in the Enterprise segment (40% of sales), the move to data center consolidation and virtualization, and the expanding role of the network; (2) in the Service Provider segment (35% of sales), the migration to next-generation networks, the advent of video on the network, and the buildout in emerging markets;

March 26, 2008

Americas: Communications Technology

and (3) in the Commercial segment (23% of sales), the adoption of turnkey solutions that also use Cisco's Advanced Technologies products, including Software-as-a-Service (SaaS). Given Cisco's solid track record of execution and core competency in making synergistic acquisitions, we expect Cisco to succeed in capitalizing on these trends to drive a midteens growth rate in the long term.

- **Margin stability.** We expect Cisco's midteens top-line growth to be mirrored on the bottom line, as gross margin pressures from incremental competition and an increased consumer mix are offset by a mix shift to higher-margin segments such as Advanced Technologies and Service Providers, as well as increased software content.

Valuation

We have raised our six-month price target to \$29 from \$28, based on 18.0X our CY2009 GAAP EPS estimate of \$1.61. Our target P/E is based on a PEG multiple of 1.2X (high end of the range in our coverage universe given the higher quality and predictability of Cisco's earnings) and a long-term EPS growth rate of 15%.

Cisco's current CY2009E multiple of 15.9X represents a 23% discount to the median growth company P/E in our group of 20.6X, and a slight discount to the S&P 500 P/E of 16.3X.

Risks

Risks include (1) a slowdown in IT spending; (2) a slowdown in carrier spending; (3) deceleration in emerging markets; (4) declines in Cisco's router and switch market share; (5) potential transition from pro forma to GAAP estimates by the Street.

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for Cisco differ from pro forma estimates primarily in that they exclude share-based compensation expense and amortization of intangibles, net of tax. We believe that consensus estimates for Cisco are currently pro forma, so we have shown changes to both our GAAP and pro forma estimates in Exhibit 22 for comparison purposes. Although the changes to our non-GAAP EPS estimates are relatively minor, we are reducing our GAAP estimates more meaningfully, primarily due to the inclusion of amortization of intangibles.

Exhibit 22: Cisco Systems

(\$ in millions)	3Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$9,755	\$39,424	\$44,448	\$50,268
GS - current	\$9,753	\$39,521	\$44,381	\$51,291
Consensus	\$9,758	\$39,464	\$44,187	\$49,223
EPS				
GS (GAAP) - previous	\$0.33	\$1.39	\$1.58	\$1.84
GS (GAAP) - current	\$0.32	\$1.34	\$1.48	\$1.74
GS (ex-ESO) - previous	\$0.36	\$1.53	\$1.73	\$1.99
GS (ex-ESO) - current	\$0.37	\$1.55	\$1.71	\$1.98
Consensus (ex-ESOs)	\$0.36	\$1.54	\$1.71	\$1.91

Source: FactSet, Goldman Sachs Research estimates.

March 26, 2008

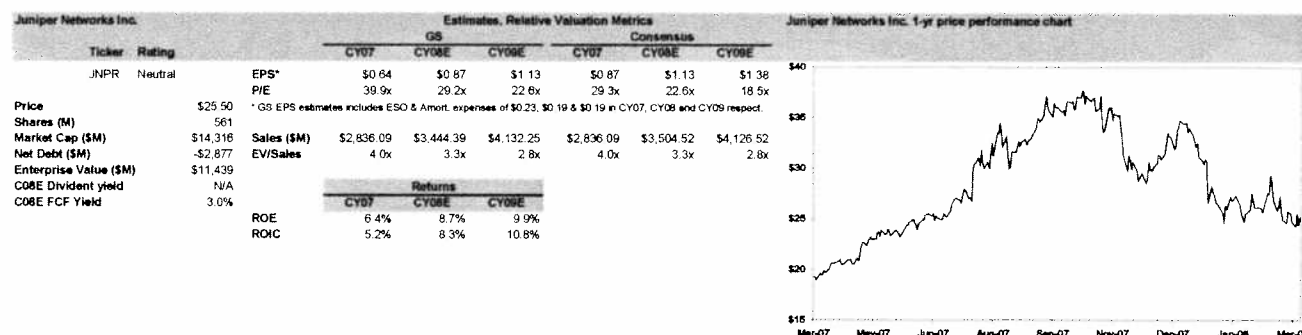
Americas: Communications Technology

Juniper Networks (JNPR, Neutral, \$25 price target)

Please also refer to our detailed Juniper report published separately today: "Assuming coverage with Neutral; solid growth, premium valuation," March 26, 2008.

Simona Jankowski has assumed primary coverage of Juniper Networks from Brantley Thompson, who is moving to the buy-side. We have moved our rating to Neutral from Buy, and lowered our price target to \$25 with a six-month horizon from \$37 with a 12-month horizon. Our price target is based on 22.0X our CY2009 GAAP EPS estimate of \$1.13, with the target P/E based on a PEG multiple of 1.1X (toward the high end of the range in our coverage universe) and a long-term EPS growth rate of 20%. Our growth expectations are based on (1) secular tailwinds such as accelerating bandwidth demand, the transition to IP-based networks, and increased network intelligence, augmented by share gains, which together should drive sales growth in the high teens, and (2) around 4% operating margin expansion through 2010 as Juniper's strong top-line growth allows for significant operating leverage.

Exhibit 23: Juniper Networks at a glance



Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- Downgrade to Neutral from Buy given macro headwinds and high bar for the stock.** Although we have a favorable view of Juniper's longer-term fundamentals, as discussed below, we believe the bar is set high for the stock in the near term, given its premium valuation and investors' high expectations for further upside to estimates. We have lowered our CY2008 estimates to slightly below consensus, as we believe the enterprise segment (29% of sales) has decelerated since the company set out its annual targets and will likely slow further; in addition, we believe that the recent pockets of weakness in the service provider segment (71% of sales) may broaden. A consistent set of input assumptions in our Cisco and Juniper models (i.e., market size, market growth, and market share) leads us to believe that the Street's estimates for Cisco have already largely priced in further deceleration, while for Juniper they have not.
- Long-term fundamentals are strong due to favorable secular trends and likely share gains.** We expect Juniper to drive top-line growth in the high teens over the next 3-5 years through a combination of market growth and share gains in its key segments: (1) in service provider routers, which drive 55% of sales, we expect Juniper to grow roughly in line with the market at around 20% going forward, as we think the company will stabilize its share after sliding in the last two years; (2) in network

March 26, 2008

Americas: Communications Technology

security, which drives 19% of sales, Juniper should drive midteens sales growth by gaining share in the fastest-growing segments of the market; and (3) while the enterprise networking market is growing only in the mid- to high single digits, Juniper will likely grow faster at a double-digit clip, as it continues to gain share off a low base in enterprise routers, which drive 7% of sales, and in addition garners 1%-2% of the enterprise switch market following its first switch introduction in early 2008.

- **Margin expansion.** In addition, we expect around 4% of GAAP operating margin expansion over the next three years, as Juniper leverages its higher revenue base and benefits from investments over the last couple of years in developing significant new router and switch platforms, as well as rationalizing acquired technologies.

Valuation

We have lowered our six-month price target to \$25, based on 22.0X our CY2009 GAAP EPS estimate of \$1.13. Our target P/E is based on a PEG multiple of 1.1X (toward the high end of the range in our coverage universe) and a long-term EPS growth rate of 20%.

Juniper's current CY2009E multiple of 22.6X represents a 10% premium to the median growth company P/E in our group of 20.6X, and a 39% premium to the S&P 500 P/E of 16.3X.

Risks

Risks include (1) a slowdown in carrier spending; (2) weakening in Juniper's competitive positioning; (3) execution in ramp of new products; (4) failure to achieve operating leverage; (5) high bar for the stock given premium multiple and high investor expectations; (6) transition from pro forma to GAAP estimates by the Street.

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for Juniper differ from pro forma estimates primarily in that they exclude share-based compensation expense and amortization of intangibles, net of tax. We believe that consensus estimates for Juniper are currently pro forma, so we have shown changes to both our GAAP and pro forma estimates below for comparison purposes.

Exhibit 24: Juniper Networks

(\$ in millions)	1Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$820	\$3,549	\$4,226	\$4,937
GS - current	\$802	\$3,444	\$4,132	\$4,951
Consensus	\$817	\$3,505	\$4,127	\$4,733
EPS				
GS - previous	\$0.21	\$0.97	\$1.29	\$1.58
GS - current	\$0.19	\$0.87	\$1.13	\$1.38
GS (ex-ESO) - previous	\$0.25	\$1.13	\$1.45	\$1.74
GS (ex-ESO) - current	\$0.23	\$1.07	\$1.32	\$1.58
Consensus (ex-ESOs)	\$0.25	\$1.13	\$1.38	\$1.51

Source: FactSet, Goldman Sachs Research estimates.


March 26, 2008

Americas: Communications Technology

Motorola, Inc. (MOT, Not Rated)

Thomas Lee is assuming coverage of Motorola from Brantley Thompson on an interim basis; coverage will transfer to Simona Jankowski after a transition period. We are Not Rated on the stock. Based on current Street estimates, Motorola is trading at 15.9X and 0.5X on 2009 P/E and EV/sales basis.

Exhibit 25: Motorola at a glance



Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Company outlook

- A number of uncertainties in 2008-2009.** Continued deterioration in Motorola's handset business is likely to continue into 2008-2009, given the lack of new compelling products on the horizon. In addition, the current lack of solid Mobile Devices leadership adds uncertainty to the timing and success of an eventual turnaround. Although the company is looking for someone with a consumer electronics background, we believe finding a strong leader who can positively influence the current culture will be challenging. Finally, the meaningful number of key personnel departures and recent appointment of a new CFO also add questions.
- Execution of software/silicon strategy.** If Motorola executes well on its transition to Texas Instruments and QUALCOMM baseband chips for its 3G handsets, it could improve the cost structure and time to market for its handset business. In addition, simplifying its software strategy (currently at five platforms) could better leverage its R&D base and ultimately, produce more compelling products on time. We note that Motorola has historically been poor in executing on these fronts.
- Stable outlook for a portion of non-handset business segments.** Currently, Motorola's Enterprise Mobility segment (made up of the former Government business and Symbol business) continues to perform well as both the Government and Symbol businesses remain well-positioned to build upon their strong performance from 2007. In Home and Networks Mobility (former Connected Home and Cellular Networks business), the Connected Home portion faces tough year-over-year comparisons (1H2007 benefited from increased set-top box sales from the FCC "707" separable security mandate) and may see some margin pressure from increased competition. In addition, the Networks business is likely to remain challenged given its subscale presence.

March 26, 2008

Americas: Communications Technology

- **Balance sheet/cash flow.** The company's net cash position has trended lower over the last year from a peak of almost \$5/share in 4Q2006 to \$2.22 at the end of 2007. With Tom Meredith's recent decision to leave the CFO position, it is unclear if the company remains as focused as before on improving its cash conversion cycle.

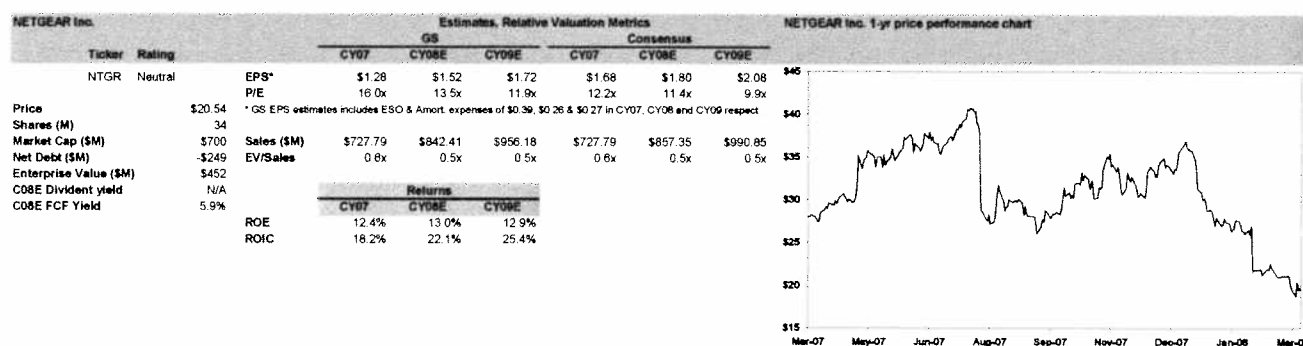
Areas of concern

Potential concerns include (1) delays in executing its cost-reduction strategy (software/silicon strategy); (2) increasing competition; and (3) continued languishing product strategy and portfolio, including the timing of compelling multimedia handsets and attractive low- and mid-end 3G handsets, and integrating the Good Technology acquisition.

Netgear (NTGR, Neutral, \$21 price target)

Thomas Lee is assuming primary coverage of Netgear from Brantley Thompson, who is moving to the buy-side. We are maintaining our Neutral rating and lowering our six-month price target to \$21 from \$27.50. Our six-month price target of \$21 is based on 12.0X our CY2009 GAAP EPS estimate of \$1.72. Our multiple is based on a target PEG multiple of 0.9X and a long-term EPS growth rate of 12%. Although current valuations are attractive (trading at 11X CY2009 GAAP EPS), we remain cautious in the near term as we head into seasonally the weakest part of the year, which we believe could be exacerbated given potential macro headwinds, particularly around consumer spending (80% of Netgear's business is in the developed regions, US and Europe, with roughly 40% of total revenues exposed to the consumer). Longer term, we believe the end-market trends remain favorable for the company and we believe the company can continue to gain share.

Exhibit 26: Netgear at a glance



Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- **Solid end-market trends, driven by low broadband penetration; market-share gains likely to continue.** Netgear provides products for wireless (WiFi) routers/access points, network interface cards or NICs, WiFi storage devices, WiFi

March 26, 2008

Americas: Communications Technology

phones, etc.) and wireline (network hubs, switches, servers, gateways, and modems) applications. We expect its end-market demand (estimated at \$2.4 billion by Infonetics) to continue to see solid growth (five-year CAGR of 10% from 2007-2012), driven in part by low broadband penetration, particularly in emerging markets (40%-50% in developed markets and 3%-4% in emerging markets). In addition, we expect Netgear to continue gain market share (WLAN revenue share at 17% in 2007 to 21% in 2009), driven in part by its continued focus in driving SMB and service provider accounts, as well as its appeal to consumers based with its sleek industrial design and easy-to-use products.

- **Solid strength in SMB and service provider segments.** Currently, 36% of Netgear's revenues are derived from the SMB segment and 22% from service providers. We believe this provides solid diversification away from the cyclical and volatile nature of the retail segment, which represents 42% of its total business. In the SMB segment, Netgear has been gaining share from companies such as 3COM and Hewlett-Packard. In the service provider channel, the company has been gaining share at key accounts such as BSKyB, Telstra in Australia, and Comcast.
- **Solid track record of financial execution; however, cautious in near term.** Overall, the company has demonstrated good execution, beating Street estimates in six out of the last eight quarters. However, in recent quarters, the company experienced a few roadblocks, missing Street estimates in two out of the last three quarters (one miss related to internal execution issues and the other to weaker-than-expected seasonal trends). While we don't necessarily believe this is a systemic trend, we note that Netgear's success is predicated largely on its ability to manage its distribution channel. If higher air freight costs persist, which were the reason for the 4Q2007 miss, the company would look to cut costs in other areas, which may be incrementally challenging given that Netgear already runs a fairly lean business.
- **High consumer exposure a concern in slowing macro environment.** With roughly 64% of its business derived from home applications and 80% of its business derived from developed markets, US and Europe (US, 38%; EMEA, 52%), we believe any material slowdown in consumer spending in either of these markets could cause Netgear to see weaker-than-expected seasonality over the coming quarters.

Valuation

Our six-month price target of \$21 is based on 12.0X our CY2009 GAAP EPS estimate of \$1.72. Our target P/E is based on a target PEG multiple of 1.0X (midpoint of the range in our coverage universe) and a long-term EPS growth rate of 12%.

Netgear's current CY2009E multiple of 11.9X represents a 42% discount to the median growth company P/E in our group of 20.6X, and a 27% discount to S&P 500 P/E of 16.3X.

Risks

Risks include (1) delays in distributing products to key markets; (2) increasing competition from established companies with deeper pockets; and (3) lack of operating leverage.

March 26, 2008

Americas: Communications Technology

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for Netgear differ from pro forma estimates primarily in that they exclude share-based compensation, net of tax. We believe that Consensus estimates for Netgear are currently pro forma, so we have shown changes to both our GAAP and pro forma estimates in Exhibit 27 for comparison purposes.

Exhibit 27: Netgear

(\$ in millions)	1Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$203	\$851	\$966	\$1,091
GS - current	\$201	\$842	\$956	\$1,052
Consensus	\$204	\$854	\$973	\$1,083
EPS				
GS (GAAP) - previous	\$0.41	\$1.71	\$1.92	\$2.17
GS (GAAP) - current	\$0.35	\$1.52	\$1.72	\$1.92
GS (ex-ESO) - previous	\$0.45	\$1.89	\$2.10	\$2.35
GS (ex-ESO) - current	\$0.41	\$1.78	\$1.99	\$2.20
Consensus (ex-ESOs)	\$0.43	\$1.80	\$2.08	NA


Source: FactSet, Goldman Sachs Research estimates.

Powerwave (PWAV, Neutral, \$2.30 price target)

Thomas Lee is assuming primary coverage of Powerwave from Brantley Thompson, who is moving to the buy-side. We are maintaining our Neutral rating and lowering our six-month price target to \$2.30 from \$4.00. Our six-month price target of \$2.30 is based on a target EV/sales multiple of 0.65X, which is near the lower end of its historical range and a slight premium to its turnaround peer group median multiple of 0.5X. Although we believe the company has made meaningful strides in its restructuring efforts, given that the company is still in a restructuring phase combined with deteriorating trends in the wireless infrastructure market, we do not believe the company should be trading a significant premium to its peer group. Longer term, we see Powerwave as a mature stock trading at a market multiple, based on its mid- to high-single-digit growth.

March 26, 2008

Americas: Communications Technology

Exhibit 28: Powerwave at a glance

Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- **We see limited downside risk in shares; however, would wait post 1Q2008 results to get more constructive on the stock.** With the stock trading at 10X our CY2009E non-GAAP EPS and 0.68X on an EV/sales basis (low end of its three-year range) we believe there is solid valuation support, particularly as we believe the company should begin to see the fruits of its restructuring efforts from the past 12-18 months. However, it may take another quarter or two before the stock begins to work, as the quality of its earnings could still be relatively poor in the near term (could see additional charges such as the goodwill impairment charges as we saw in 4Q2007). As we look beyond 1Q2008, however, we think there could be upside in the stock, as we believe there is reasonable probability that the company reports full-year 2008 revenues at or above the midpoint of its guided range of \$860-\$900 million due to a 3G network ramp at AT&T.
- **Still more room to go on cost-cutting side (halfway through restructuring phase).** In our view, the company is in the middle stages of its restructuring, and we see opportunities for lower costs (both COGS and opex) on top of a stable revenue base to drive steady margin expansion going forward. We are forecasting GAAP operating margins of -7%, -3%, and 1% in 2008, 2009, and 2010 and non-GAAP operating margins of 0%, 2%, and 3%. From a cost-cutting perspective, the company has done a good job reducing its overall opex level from \$62 million in 1Q2007 to \$54 million in 4Q2007. We believe there is still much more room to go on this front, both in research and development and sales and marketing. Management is currently targeting a non-GAAP opex level of \$45 million sometime in 2H2008, with lower levels in 2009 and beyond. In addition, we believe the company has additional opportunity to lower its cost base and improve gross margins, for instance by selling additional manufacturing plants, following the two it sold in 2007. With the recent sale of its Hungary manufacturing plant to Sanmina, we believe the company will see improving gross margins in coming quarters.
- **Longer term, company is likely to be a mid- to high-single-digit earnings grower.** Although we are forecasting the wireless infrastructure market to see virtually zero growth over the next three years, we believe that Powerwave is likely to see higher growth based on (1) potential market-share gains given CommScope's recent acquisition of Andrew, Powerwave's top competitor, as CommScope/Andrew is likely to exit certain businesses/accounts, which are non-strategic, (2) a niche market for Powerwave as operators will need to maintain

March 26, 2008

Americas: Communications Technology

current networks by upgrading antennas, power amplifiers, and other components that large system OEM equipment vendors would not necessarily want to support.

Valuation

Our six-month price target of \$2.30 is based on 0.65X our CY2009 sales estimate of \$949 million, which is near the lower end of its historical range and roughly in line with its turnaround peer group median multiple of 0.5X. Although we believe the company has made meaningful strides in its restructuring efforts, given that the company is still in a restructuring phase and is in a broader end market (wireless infrastructure) that is deteriorating, we do not believe Powerwave shares should trading a significant premium to its peer group.

Currently, the stock is trading at a non-GAAP CY2009E P/E of 9.4X and EV/sales of 0.68X.

Risks

Risks include (1) continued deterioration in wireless infrastructure market; (2) weak spending patterns at large customers; and (3) pricing pressure

Estimate changes

We are transitioning to GAAP EPS estimates as our primary earnings metric. Our GAAP estimates for Powerwave (which we believe are most comparable to the Street) includes share-based compensation but excludes amortization of intangibles. We exclude amortization of intangibles from our estimates because a large percentage of its current intangibles are largely a result from one acquisition (Filtronics in 4Q2006).

Exhibit 29: Powerwave

(\$ in millions)	1Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$205	\$880	\$949	\$997
GS - current	\$205	\$880	\$949	\$997
Consensus	\$207	\$883	\$963	NA
EPS				
GS (GAAP) - previous	(\$0.02)	\$0.07	\$0.27	\$0.32
GS (GAAP) - current	(\$0.03)	\$0.02	\$0.24	\$0.32
Consensus (GAAP)	(\$0.03)	\$0.08	\$0.27	NA

Source: FactSet, Goldman Sachs Research estimates.

QUALCOMM (QCOM, Buy, \$44 price target)

Thomas Lee is assuming coverage of QUALCOMM from Brantley Thompson on an interim basis and maintaining our Buy rating; coverage will transfer to Simona Jankowski after a transition period. We are moving to a six-month target of \$44 (vs. prior 12-month target of \$58), which is based on 20.7X our CY2009

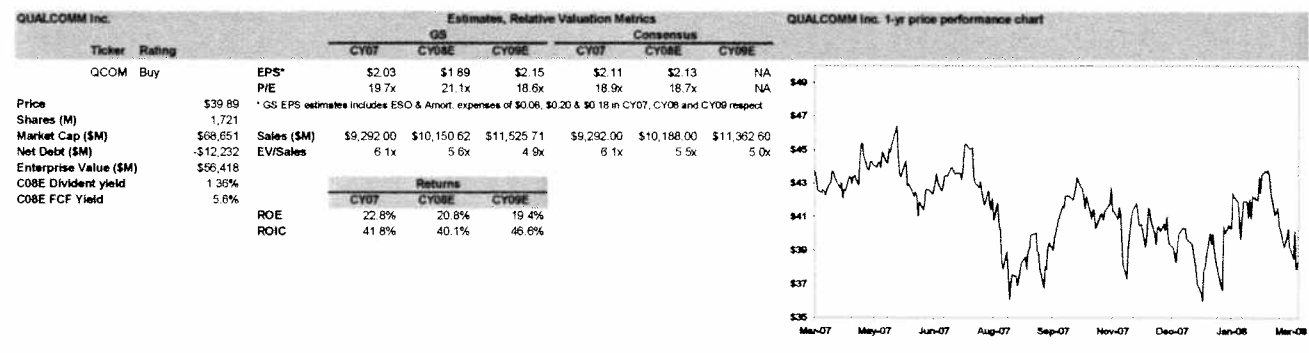
March 26, 2008

Americas: Communications Technology

GAAP EPS estimate of \$2.15. Our multiple is based on a target PEG multiple of 1.1X and a long-term EPS growth rate of 18%.

Our positive view on QUALCOMM is based on (1) solid CDMA and WCDMA handset unit growth (two-year CAGR from CY2007 to 2009 of 6% and 40%, respectively), which should drive solid licensing revenues over the next several years; for context, we estimate CDMA represents 43% and WCDMA 48% of QUALCOMM's revenues; (2) steady 3G chipset market share gains due to QUALCOMM's technology lead over rivals such as Freescale and Texas Instruments; (3) legal resolution between Nokia and QUALCOMM, which we estimate will drive \$0.20 to \$0.25 earnings upside in the first full fiscal year of resolution (assuming Nokia pays a 2.0% to 2.5% royalty rate, the midpoint of the prior rate).

Exhibit 30: QUALCOMM at a glance



Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- Continued strong growth in WCDMA market with further upside to estimates.** We are forecasting the WCDMA handset market to grow at a five-year unit CAGR of 27% from 2007 to 2012, compared to the overall handset market of 9%. We expect the mix of WCDMA handsets, currently 16% of the market, to increase to 34% in 2012. By contrast, the GSM market, which currently represents 65% of the total, should decline to 50%. We believe that our WCDMA estimates could prove to be conservative depending on how 3G penetration progresses in emerging markets such as India, Eastern Europe, and North Asia. As a reference point, in November 2007, QUALCOMM guided the CY2008 WCDMA handset market to be 284 million units (+56% year over year).
- Healthy CDMA market expected over next several years.** We believe the CDMA market will continue to see reasonable growth in the low- to mid-single digits from 2007 to 2009. We expect to see a reacceleration in CDMA demand in developing regions such as Africa, India, Indonesia, and Malaysia, driven in part by wider availability of lower-cost CDMA handsets and increasing usage of data. Furthermore, we expect to see a healthy replacement market in developed countries such as the US, Japan, and Korea, as operators continue to migrate their existing 1X subscribers to EV-DO networks.
- Extending its leadership in 3G baseband market.** We believe QUALCOMM is poised for continued share gains in the WCDMA chipset market, both at the high end

March 26, 2008

Americas: Communications Technology

with HSPA and at the low end with its single-chip WCDMA solution. We expect to see an aggressive 3G push in 2008 from QUALCOMM's key customers such as HTC, LG, Motorola, and Samsung. Despite increasing chatter of increased 3G baseband competition, we believe QUALCOMM will continue to grow faster than the market as we believe it has the largest and broadest 3G chipset portfolio in the marketplace. We believe QUALCOMM is clearly benefiting from the large R&D investment in 3G chipset development it has made over the past several years, which will make it increasingly difficult for the competition to close the gap in the near to medium term.

- **Current valuations reflect an overly negative outcome regarding current Nokia licensing situation.** In addition, with shares pricing in an overly negative outcome regarding the current Nokia license renegotiations, meeting legal timetables throughout 2008 should help provide relief to shares. In our sum-of-the-parts analysis, current valuation reflects an unrealistic scenario that Nokia never pays QUALCOMM a royalty payment and the royalty rate for all other licensees is reduced to 2.0% in FY2008-FY2012 and roughly 1.0% in FY2013 and beyond, from around 4.0% currently. We find this unlikely, given that only two WCDMA subscriber licensees are set to expire before November 2017. We expect to see resolution to some of the key legal issues during 2008, thus alleviating a key overhang on the shares.

Valuation

We are moving to a six-month price target of \$44 (vs. prior 12-month target of \$58), which is based on 20.7X our CY2009 GAAP EPS estimate of \$2.15. Our target P/E is based on a target PEG multiple of 1.0X (midpoint of range in our coverage universe) and a long-term EPS growth rate of 18%.

QUALCOMM's current CY2009E multiple of 18.6X represents a 10% discount to the median growth company P/E in our group of 20.6X, and a 14% premium to S&P 500 P/E of 16.3X.

Risks

Risks include (1) changes to the royalty rate, (2) increased WCDMA chipset competition, (3) CDMA/WCDMA handset market weakness, (4) a further deterioration in the macro environment.

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for QUALCOMM differ from pro forma estimates primarily in that they exclude share-based compensation expense, net of tax. We believe that consensus estimates for QUALCOMM are currently pro forma, so we have shown changes to both our GAAP and pro forma estimates in Exhibit 31 for comparison purposes.

March 26, 2008

Americas: Communications Technology

Exhibit 31: QUALCOMM estimate changes

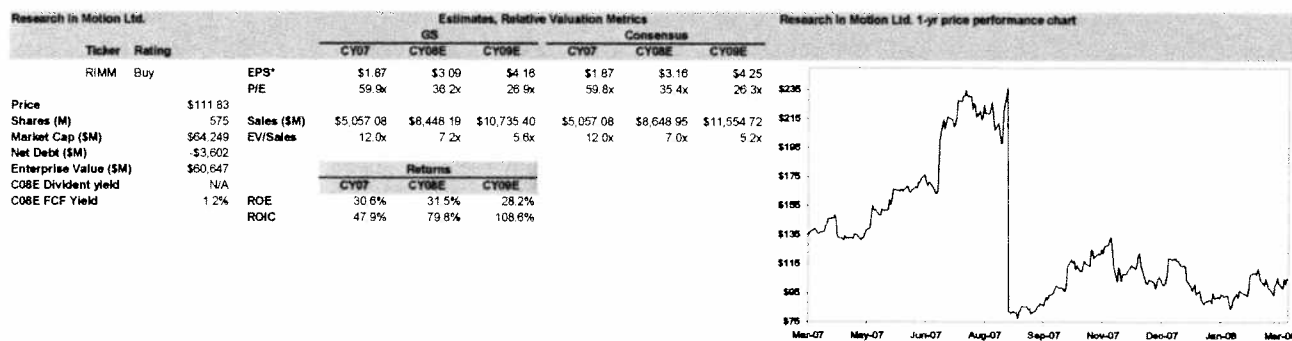
(\$ in millions)	2Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$2,534	\$10,012	\$11,356	\$12,827
GS - current	\$2,516	\$9,915	\$11,168	\$12,399
Consensus	\$2,503	\$9,946	\$11,153	\$12,910
EPS				
GS (GAAP) - previous	\$0.47	\$1.88	\$2.12	\$2.39
GS (GAAP) - current	\$0.47	\$1.86	\$2.09	\$2.30
GS (ex-ESO) - previous	\$0.52	\$2.07	\$2.31	\$2.58
GS (ex-ESO) - current	\$0.52	\$2.07	\$2.28	\$2.46
Consensus (ex-ESO)	\$0.52	\$2.10	\$2.39	\$2.80

Source: FactSet, Goldman Sachs Research estimates.

Research in Motion (RIMM, Buy, \$120 price target)

Thomas Lee is assuming coverage of Research in Motion from Brantley Thompson on an interim basis, and maintaining our Buy rating; coverage will transfer to Simona Jankowski after a transition period. Our six-month price target of \$120 is based on 28.7X our CY2009 GAAP EPS estimate of \$4.16, and is down slightly from \$125 previously. Our target P/E multiple is based on a target PEG of 1.2X (high-end of range in our coverage universe) and long-term EPS growth rate of 23%.

Our robust growth expectations are based on (1) RIM's leadership in wireless email (enterprise and consumer), a market that in the US alone is expected to grow at three-year CAGR of 44% from 2007 to 2010 and that is still less than 5% penetrated, and (2) the company's solid execution track record, which bolsters our confidence in its ability to expand its market share through penetration of overseas markets and the "prosumer" segment, beyond its traditional beachhead in the North America enterprise market.

Exhibit 32: Research in Motion at a glance

Source: Goldman Sachs Research estimates, Thomson ONE, FactSet.

Investment view

- **Strong growth in the wireless e-mail market.** RIM's primary end market is wireless e-mail, which is growing in excess of 40% annually driven by the trend toward an increasingly mobile workforce and richer consumer handheld experiences. Our US Telecom Services analyst Jason Armstrong is forecasting US wireless data revenues to increase to \$58 billion in 2010 from \$24 billion in 2007 (three-year CAGR of 35%), of which BlackBerry service revenues are expected to represent 23% in 2010 from 17% in 2007 (three-year CAGR of 44%). Currently, out of 600 million corporate e-mail accounts, only 8 million are wirelessly connected through RIM's devices, and RIM remains far ahead of competitors such as Microsoft, Nokia/Intellisync, and Motorola/Good. This means that RIM's penetration is less than 2%, and the overall market penetration is well under 5%. The greenfield opportunity is even more significant on the consumer side, given 3 billion mobile subscribers and only 4 million RIM consumer wireless e-mail subscribers.
- **Expanding to new customer and geographic segments.** RIM has established a formidable beachhead in certain verticals of the enterprise e-mail market, such as the financial community, government, and legal. Based on the its solid track record of execution and low market penetration, we believe RIM has significant growth opportunities going forward through (1) expanding its distribution channels overseas and (2) moving beyond its traditional enterprise customer base to the consumer/prosumer segment. Consequently, we believe RIM can gain share in the fast-growing smartphone segment as we estimate its smartphone market share will reach 15% by 2011 from 10% at the end of 2007. Note that only around one-third of RIM's subscriber base is currently outside the US. In 2008, we expect growth from additional retail distributors in the US and Western Europe, continued Latin America growth, a significant initiative in Eastern Europe and Russia, the 8700 launch in China in early 2008, and a broader push by DoCoMo in Japan. In addition, the 2008 product pipeline is compelling, with new 3G devices, the 9000 series, CDMA Curve, more Wi-Fi embedded devices, and BlackBerry Unite in mid-2008.
- **Operating margin expansion in FY2008-FY2010.** We are forecasting operating margin expansion from 27% in FY2007 (February) to 31% in FY2010, as we expect RIM's higher-margin services business segment to increase to 22% of total revenues from 18% in FY2007. This is largely driven by strong subscriber additions, as we expect RIM's total BlackBerry subscriber base to increase to 41.5 million at the end of FY2010 from 8.2 million at the end of FY2007 (3-year CAGR of 70%). We expect the mix of BlackBerry Internet Subscribers (BIS), which are a mix of SMB and consumer subscribers, to increase to 43% in FY2010 from 27% at the end of FY2007.

Valuation

We are lowering our six-month price target slightly to \$120 from \$125, based on 28.7X our CY2009 GAAP EPS estimate of \$4.16. Our target P/E is based on a PEG multiple of 1.2X (high-end of range in our coverage universe) and a long-term EPS growth rate of 23%

RIM's current CY2009 multiple of 26.9X represents a 31% premium to the median growth company P/E in our group of 20.6X, and a 65% premium to S&P 500 P/E of 16.3X.

Risks

Risks include (1) lower enterprise IT spending, (2) increasing competition, (3) increased investment may reduce operating leverage, (4) network operating center (NOC) scalability.

March 26, 2008

Americas: Communications Technology

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for RIM differ from pro forma estimates primarily in that they exclude share-based compensation expense, net of tax.

Exhibit 33: Research in Motion

(\$ in millions)	4Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$1,838	\$5,964	\$8,784	\$11,454
GS - current	\$1,878	\$6,005	\$9,047	\$11,526
Consensus	\$1,846	\$5,941	\$9,083	\$11,895
EPS				
GS (GAAP) - previous	\$0.69	\$2.20	\$3.24	\$4.50
GS (GAAP) - current	\$0.70	\$2.24	\$3.34	\$4.50
Consensus (GAAP)	\$0.70	\$2.24	\$3.44	\$4.64

Source: Factset, Goldman Sachs Research estimates.

Starent Networks (STAR, Neutral, \$14 price target)

Thomas Lee is assuming primary coverage of Starent Networks from Brantley Thompson, who is moving to the buy-side. We are downgrading Starent to Neutral from Buy and reducing our six-month price target to \$14 from \$23. Our six-month price target of \$14 is based on 29.3X our CY2009 GAAP EPS estimate of \$0.48. Our multiple is based on a target PEG multiple of 1.0X and a long-term EPS growth rate of 28%.

We continue to believe the near-term fundamentals remain strong for Starent, due to its exposure to wireless data growth which is still in its early stages, a high-quality customer base, and a strong competitive position in the marketplace. However, given the current macro headwinds and investors' valuation sensitivity, we have a stronger bias toward large-cap stocks such as Cisco, QUALCOMM, and Research in Motion that have better geographic and vertical diversification. With that said, Starent is our favorite Neutral-rated stock.

March 26, 2008

Americas: Communications Technology

positioned to win large tier-1 accounts, we believe tier-2 operators may not necessarily choose best-in-breed as price and other considerations may play a more important role, which could favor larger vendors that offer a larger suite of products (full-system offering). For instance, T-Mobile selected Huawei as a supplier for next-generation wireless gateways in several European regions, which we believe was largely due to Huawei's wireline offering rather than purely price or performance considerations on the wireless equipment side.

- **Rich valuation.** Although we view Starent's long-term business outlook as promising, we see limited near-term upside given that the shares are currently trading at 30X our CY2009 GAAP EPS estimate (22X on a non-GAAP basis) versus our small-cap group median of 18X. Although Starent's business may be more resilient in an economic downturn, it is not likely to see multiple expansion in an environment when investors have increased sensitivity to high multiples and risk aversion to small-cap stocks.

Valuation

We are lowering our six-month price target to \$14 (previously \$23), which is based on 29.3X our CY2009 GAAP EPS estimate of \$0.48. Our target P/E is based on a PEG multiple of 1.0X (midpoint of range in our coverage universe) and a long-term EPS growth rate of 28%.

Starent's current CY2009E multiple of 30.6X represents a 48% premium to the median growth company P/E in our group of 20.6X, and a 88% premium to S&P 500 P/E of 16.3X.

Risks

Risks include (1) increased competition; (2) high customer concentration; (3) lumpy order patterns.

Estimate changes

We are transitioning to GAAP EPS estimates as our primary valuation metric. GAAP estimates for Starent differ from pro forma estimates primarily in that they exclude share-based compensation expense, net of tax.

Exhibit 35: Starent

(\$ in millions)	1Q08E	2008E	2009E	2010E
Revenues				
GS - previous	\$51	\$219	\$296	\$381
GS - current	\$51	\$219	\$296	\$381
Consensus	\$48	\$212	\$281	\$381
EPS				
GS (GAAP) - previous	\$0.05	\$0.45	\$0.50	\$0.75
GS (GAAP) - current	\$0.04	\$0.40	\$0.48	\$0.65
Consensus (GAAP)	\$0.05	\$0.39	\$0.47	\$0.75

Source: FactSet, Goldman Sachs Research estimates.

March 26, 2008

Americas: Communications Technology

Removing Starent (STAR) from the Americas Buy List

We are removing Starent Networks from the Americas Buy List and downgrading it to Neutral given our current bias toward large-cap stocks as we see limited upside to shares in the near term given current macro headwinds and investors' valuation sensitivity. With that said, Starent is our favorite Neutral-rated stock. Since being added to the Buy List on July 16, 2007, STAR shares are down 1.4% versus down 12.7% for the S&P 500.

March 26, 2008

Americas: Communications Technology

Exhibit 36: Starent (STAR) share price performance versus peer group, part I
prices as of the close of March 25, 2008

Company	Ticker	Primary analyst	Price currency	Price as of 03/25/08	Price performance since 07/16/07	3 month price performance	6 month price performance	12 month price performance
Americas Technology Peer Group								
Starent Networks, Corp.	STAR	Brantley Thompson	\$	14.81	-1.4%	-25.7%	-28.3%	NA
3PAR INC	PAR	David C. Bailey	\$	6.88	NA	-51.8%	NA	NA
Accenture, Ltd.	ACN	Julio C. Quinteros, Jr.	\$	35.38	-16.8%	-7.0%	-7.0%	-2.8%
Acme Packet, Inc.	APKT	Brantley Thompson	\$	7.77	-37.4%	-40.6%	-49.6%	-47.9%
Adobe Systems Inc.	ADBE	Sasa Zorovic, Ph.D	\$	36.54	-11.0%	-14.7%	-14.5%	-14.6%
Advanced Energy Industries, Inc.	AEIS	James Covello	\$	13.61	-39.3%	1.6%	-12.8%	-36.1%
Advanced Micro Devices, Inc.	AMD	James Covello	\$	6.30	-59.9%	-18.9%	-54.4%	-53.3%
Affiliated Computer Services, Inc.	ACS	Julio C. Quinteros, Jr.	\$	51.05	-9.5%	11.6%	3.2%	-14.0%
Akamai Technologies, Inc.	AKAM	Derek Bingham	\$	32.59	-33.4%	-10.6%	8.1%	-36.6%
Altera Corp.	ALTR	James Schneider, Ph.D.	\$	18.86	-22.4%	-3.9%	-24.7%	-12.1%
Amdocs Limited	DOX	Elizabeth W. Gausam, CFA	\$	29.76	-24.2%	-12.3%	-15.4%	-17.3%
Analog Devices, Inc.	ADI	Simona Jankowski, CFA	\$	29.51	-24.8%	-7.7%	-20.4%	-17.0%
Apple Inc.	AAPL	David C. Bailey	\$	140.98	2.1%	-29.1%	-8.0%	50.7%
Applied Materials, Inc.	AMAT	James Covello	\$	20.66	0.2%	13.8%	-1.7%	10.0%
Applied Micro Circuits	AMCC	James Schneider, Ph.D.	\$	7.60	-31.7%	-20.6%	-40.1%	-45.4%
Aruba Networks, Inc.	ARUN	Brantley Thompson	\$	5.95	-73.5%	-60.7%	-68.3%	NA
ATMI, Inc.	ATMI	James Covello	\$	28.08	-12.5%	-14.3%	-5.6%	-14.8%
Autodesk Inc.	ADSK	Sasa Zorovic, Ph.D	\$	33.12	-29.6%	-35.1%	-33.6%	-13.7%
Automatic Data Processing Inc.	ADP	Elizabeth W. Gausam, CFA	\$	42.43	-13.6%	-7.3%	-8.7%	-6.9%
Axcelsis Technologies, Inc.	ACLS	James Covello	\$	5.66	-8.9%	15.0%	11.6%	-27.0%
BEA Systems, Inc.	BEAS	Sarah Friar	\$	19.14	42.3%	20.0%	44.2%	62.9%
BearingPoint, Inc.	BE	Julio C. Quinteros, Jr.	\$	1.79	-76.1%	-39.3%	-58.1%	-76.5%
BMC Software, Inc.	BMC	Derek Bingham	\$	33.62	11.9%	-8.1%	6.1%	4.9%
Broadcom Corporation	BRCM	James Schneider, Ph.D.	\$	18.83	-41.9%	-29.4%	-49.0%	-41.7%
Brocade Communications Systems	BRCD	Min Park	\$	7.55	-2.3%	1.6%	-7.5%	-24.7%
Brooks Automation, Inc.	BRKS	James Covello	\$	10.26	-43.0%	-24.1%	-26.6%	-39.6%
CA, Inc.	CA	Sarah Friar	\$	23.18	-11.5%	-10.5%	-10.1%	-13.4%
Check Point Software Tech.	CHKP	Sarah Friar	\$	24.06	3.8%	7.2%	-3.1%	4.2%
CIENA Corporation	CIEN	Brantley Thompson	\$	31.88	-17.6%	-7.8%	-15.1%	19.4%
Cisco Systems, Inc.	CSCO	Brantley Thompson	\$	25.75	-13.9%	-10.3%	-20.6%	-1.7%
Citrix Systems Inc.	CTXS	Sarah Friar	\$	32.62	-6.2%	-14.2%	-17.1%	0.4%
Cognizant Technology Solutions	CTSH	Julio C. Quinteros, Jr.	\$	30.56	-28.7%	-14.5%	-20.9%	-34.2%
CommVault Systems, Inc.	CVLT	Derek Bingham	\$	14.39	-18.4%	-35.0%	-22.1%	-10.5%
Computer Sciences Corp.	CSC	Julio C. Quinteros, Jr.	\$	41.59	-31.0%	-20.2%	-28.0%	-20.6%
Converse Technology, Inc.	CMVT	Brantley Thompson	\$	15.00	-32.1%	-11.3%	-26.5%	-28.3%
Convergys Corporation	CVG	Elizabeth W. Gausam, CFA	\$	15.32	-36.7%	-10.3%	-10.4%	-40.5%
Corning Incorporated	GLW	Brantley Thompson	\$	25.14	-6.6%	4.8%	4.4%	9.4%
Credence Systems Corp.	CMOS	James Covello	\$	1.50	-60.6%	-45.1%	-51.0%	-56.0%
CSG Systems International, Inc.	CSGS	Elizabeth W. Gausam, CFA	\$	11.31	-57.8%	-27.0%	-47.1%	-54.6%
Data Domain, Inc.	DDUP	David C. Bailey	\$	23.76	-15.1%	-8.3%	-22.7%	NA
Dell Inc.	DELL	David C. Bailey	\$	20.42	-29.5%	-18.8%	-27.5%	-10.6%
Digital River, Inc.	DRIV	Sasa Zorovic, Ph.D	\$	31.84	-33.4%	-6.7%	-28.8%	-42.9%
Electronic Data Systems	EDS	Julio C. Quinteros, Jr.	\$	17.22	-39.3%	-19.8%	-20.8%	-38.1%
EMC Corporation	EMC	David C. Bailey	\$	14.94	-21.9%	-21.6%	-28.6%	9.1%
Emulex Corp.	ELX	Min Park	\$	16.60	-27.6%	-2.7%	-11.7%	-9.3%
Entegris, Inc.	ENTG	James Covello	\$	7.31	-37.0%	-21.1%	-18.9%	-33.1%
ExdService Holdings, Inc.	EXLS	Julio C. Quinteros, Jr.	\$	23.12	28.4%	-0.4%	7.0%	1.9%
Fidelity National Information Svcs.	FIS	Julio C. Quinteros, Jr.	\$	40.14	-30.4%	-6.2%	-10.0%	-8.5%
Fiserv, Inc.	FISV	Julio C. Quinteros, Jr.	\$	49.56	-12.7%	-10.0%	-1.8%	-7.3%
FormFactor, Inc.	FORM	James Covello	\$	18.67	-50.2%	-44.5%	-59.2%	-59.3%
Genpact Ltd.	G	Julio C. Quinteros, Jr.	\$	13.48	NA	-15.5%	-1.6%	NA
Global Cash Access Holdings	GCA	Elizabeth W. Gausam, CFA	\$	6.03	-62.9%	-11.8%	-44.1%	-62.9%
Global Payments Inc.	GPNI	Elizabeth W. Gausam, CFA	\$	39.36	-0.8%	-13.1%	-8.3%	2.4%
Heartland Payment Systems, Inc.	HPY	Elizabeth W. Gausam, CFA	\$	21.51	-27.7%	-22.2%	-16.4%	-10.6%
Hewitt Associates, Inc.	HEW	Elizabeth W. Gausam, CFA	\$	40.00	23.6%	3.7%	14.4%	35.7%
Hewlett-Packard Co.	HPQ	David C. Bailey	\$	48.26	1.7%	-7.8%	-5.3%	19.4%
Infinera Corp.	INFN	Brantley Thompson	\$	13.07	-45.1%	-10.3%	-30.9%	NA
Informatica Corp.	INFA	Sasa Zorovic, Ph.D	\$	17.57	17.0%	-2.5%	13.2%	25.9%
Infosys Technologies Ltd. (ADR)	INFY	Julio C. Quinteros, Jr.	\$	37.40	-28.3%	-19.2%	-19.1%	-29.1%
Ingram Micro Inc.	IM	Min Park	\$	16.16	-25.5%	-14.7%	-16.4%	-18.0%
Intel Corp.	INTC	James Covello	\$	22.27	-14.2%	-18.5%	-14.0%	15.6%
International Business Machines	IBM	David C. Bailey	\$	117.97	7.6%	5.7%	1.3%	24.1%
International Rectifier Corp.	IRF	Simona Jankowski, CFA	\$	22.10	-43.2%	-35.3%	-31.0%	-44.1%
Intersil Corp.	ISIL	Simona Jankowski, CFA	\$	27.31	-18.3%	8.8%	-17.7%	3.8%
Intevac, Inc.	IVAC	Hongyu Cai, CFA	\$	13.82	-32.4%	-10.4%	-8.2%	-49.0%
Isilon Systems, Inc.	ISLN	David C. Bailey	\$	5.18	-67.4%	2.6%	-37.3%	-70.5%
JDS Uniphase Corp.	JDSU	Brantley Thompson	\$	14.11	-9.4%	2.0%	0.1%	-8.8%
Juniper Networks, Inc.	JNPR	Brantley Thompson	\$	26.58	-1.4%	-22.7%	-25.6%	40.3%
Kenexa Corporation	KNXA	Sasa Zorovic, Ph.D	\$	19.35	-47.7%	-6.5%	-37.4%	-40.4%
KLA-Tencor	KLAC	James Covello	\$	37.40	-34.0%	-25.1%	-34.1%	-31.2%
Lam Research Corp.	LRCX	James Covello	\$	38.25	-29.8%	-14.8%	-29.4%	-19.2%
Lattice Semiconductor Corp.	LSCC	James Schneider, Ph.D.	\$	2.80	-50.8%	-15.7%	-41.4%	-54.5%
Leadis Technology	LDIS	James Schneider, Ph.D.	\$	1.89	-46.0%	-27.9%	-42.4%	-53.0%
Lexmark International Group	LXK	Min Park	\$	31.66	-31.2%	-10.7%	-16.9%	-47.1%

Note: This table shows movement in absolute share price and not total shareholder return.
Results presented should not and cannot be viewed as an indicator of future performance.

Source: FactSet, Quantum database.

March 26, 2008

Americas: Communications Technology

Exhibit 36, cont'd: Starent (STAR) share price performance versus peer group, part II
prices as of the close of March 25, 2008

Limelight Networks, Inc.	LLNW	Derek Bingham	\$	3.34	-83.3%	-54.2%	-64.5%	NA
Linear Technology Corp.	LLTC	Simona Jankowski, CFA	\$	31.67	-16.2%	-1.3%	-8.9%	-4.2%
Lionbridge Technologies, Inc.	LIOX	Christopher Agnew, CFA	\$	3.70	-29.4%	1.4%	0.0%	-30.1%
LSI Corp.	LSI	James Schneider, Ph.D.	\$	5.42	-30.2%	-2.9%	-23.1%	-47.2%
Macrovision Corp.	MVSN	Sasa Zorovic, Ph.D.	\$	14.24	-44.4%	-25.0%	-39.8%	-42.6%
Marvell Technology Group Ltd.	MRVL	James Schneider, Ph.D.	\$	11.31	-39.3%	-20.0%	-29.3%	-37.9%
Mastercard Inc.	MA	Elizabeth W. Gausam, CFA	\$	224.98	32.6%	6.1%	57.3%	104.8%
Maxim Integrated Products	MXIM	Simona Jankowski, CFA	\$	20.30	-41.8%	-20.5%	-28.0%	-35.4%
McAfee, Inc.	MFE	Sarah Friar	\$	34.94	-3.6%	-9.2%	-0.7%	16.5%
Micrel, Inc.	MCRL	Simona Jankowski, CFA	\$	9.38	-31.4%	7.3%	-15.0%	-18.4%
Microchip Technology Inc.	MCHP	Simona Jankowski, CFA	\$	34.56	-9.2%	5.7%	-5.4%	-5.3%
Micron Technology Inc.	MU	James Covello	\$	5.79	-58.6%	-23.3%	-44.4%	-50.0%
Microsoft Corp.	MSFT	Sarah Friar	\$	29.14	-3.0%	-20.3%	-1.4%	4.0%
MKS Instruments, Inc.	MKSI	James Covello	\$	22.08	-15.7%	12.7%	11.8%	-14.7%
MoneyGram International, Inc.	MGI	Elizabeth W. Gausam, CFA	\$	2.33	-92.3%	-85.3%	-89.0%	-91.8%
Monolithic Power Systems	MPWR	Simona Jankowski, CFA	\$	18.05	-1.0%	-18.7%	-28.1%	37.5%
Motorola, Inc.	MOT	Brantley Thompson	\$	9.76	-46.3%	-40.3%	-47.6%	-45.0%
National Semiconductor Corp.	NSM	Simona Jankowski, CFA	\$	18.91	-34.9%	-19.8%	-29.3%	-23.5%
Netgear, Inc.	NTGR	Brantley Thompson	\$	20.51	-48.1%	-43.4%	-28.9%	-27.2%
Network Appliance Inc.	NTAP	Min Park	\$	20.95	-28.9%	-18.7%	-22.1%	-45.1%
NeuStar, Inc.	NSR	Elizabeth W. Gausam, CFA	\$	26.29	-13.2%	-6.3%	-20.4%	-11.7%
Nortel Networks Corporation	NT	Brantley Thompson	\$	6.85	-70.6%	-56.2%	-56.8%	-71.5%
Novellus Systems Inc.	NVLS	James Covello	\$	22.24	-25.1%	-20.0%	-18.2%	-31.9%
Nuance Communications, Inc.	NUAN	Derek Bingham	\$	18.19	2.1%	-5.6%	-6.3%	13.8%
Nvidia Corp.	NVDA	Simona Jankowski, CFA	\$	20.32	-34.7%	-43.2%	-43.3%	0.6%
Opnext, Inc.	OPXT	Brantley Thompson	\$	5.86	-59.6%	-37.4%	-46.6%	-64.4%
Oracle Corp.	ORCL	Sarah Friar	\$	21.08	4.4%	-7.4%	-3.9%	15.6%
Parametric Technology Corp.	PMTC	Sasa Zorovic, Ph.D.	\$	16.40	-6.3%	-11.2%	-2.4%	-16.2%
Patni Computer Systems Ltd. (ADR)	PTI	Julio C. Quinteros, Jr.	\$	11.50	-56.5%	-33.4%	-49.0%	-50.4%
Paychex, Inc.	PAYX	Elizabeth W. Gausam, CFA	\$	33.29	-24.2%	-11.2%	-23.2%	-18.3%
Pitney Bowes Inc.	PBI	Julio C. Quinteros, Jr.	\$	36.17	-22.3%	-5.6%	-19.8%	-21.5%
PMC-Sierra, Inc.	PMCS	James Schneider, Ph.D.	\$	5.81	-29.2%	-17.2%	-29.7%	-12.6%
Powerwave Technologies, Inc.	PWAV	Brantley Thompson	\$	2.62	-61.1%	-41.8%	-57.4%	-52.5%
QLogic Corp.	QLGC	Min Park	\$	15.25	-11.4%	3.0%	18.6%	-11.5%
QUALCOMM, Inc.	QCOM	Brantley Thompson	\$	40.80	-9.4%	1.2%	-1.8%	-4.2%
Quest Software, Inc.	QSFT	Derek Bingham	\$	13.41	-17.0%	-27.8%	-19.9%	-18.9%
Red Hat, Inc.	RHT	Sarah Friar	\$	18.20	-15.5%	-13.9%	-3.7%	-24.9%
Research In Motion Ltd.	RIMM	Brantley Thompson	\$	115.95	53.5%	-1.7%	19.8%	153.4%
Research In Motion Ltd.	RIM.TO	Brantley Thompson	C\$	118.31	50.3%	1.6%	21.6%	122.6%
RightNow Technologies, Inc.	RNOW	Sasa Zorovic, Ph.D.	\$	11.44	-30.4%	-29.7%	-26.7%	-32.3%
Riverbed Technology, Inc.	RVBD	Brantley Thompson	\$	16.72	-64.3%	-36.9%	-61.9%	-39.3%
Salesforce.com, Inc.	CRM	Sasa Zorovic, Ph.D.	\$	60.96	38.1%	-4.7%	22.1%	40.8%
SanDisk Corporation	SNDK	James Covello	\$	21.76	-59.8%	-38.5%	-58.4%	-51.5%
SAP (ADR)	SAP	Sarah Friar	\$	51.42	-1.9%	-1.0%	-12.0%	12.8%
Sapient	SAPE	Julio C. Quinteros, Jr.	\$	7.67	-0.1%	-15.2%	15.2%	13.1%
Satyam Computer Svcs. Ltd. (ADR)	SAY	Julio C. Quinteros, Jr.	\$	24.16	-8.7%	-13.5%	0.5%	5.2%
Seagate Technology	STX	David C. Bailey	\$	21.63	-5.0%	-17.0%	-18.6%	-13.6%
Secure Computing Corp.	SCUR	Sarah Friar	\$	7.62	-6.7%	-22.2%	-18.2%	-8.6%
Sigma Designs, Inc.	SIGM	James Schneider, Ph.D.	\$	23.93	-21.7%	-58.0%	-52.0%	-22.5%
Silicon Image, Inc.	SIMG	James Schneider, Ph.D.	\$	4.94	-40.2%	-0.6%	-11.9%	-44.8%
SIRF Technology Holdings, Inc.	SIRF	James Schneider, Ph.D.	\$	4.93	-78.8%	-80.3%	-75.1%	-82.3%
Sonus Networks, Inc.	SONS	Brantley Thompson	\$	3.48	-58.3%	-42.5%	-39.4%	-56.7%
SuccessFactors, Inc.	SFSF	Sarah Friar	\$	10.60	NA	-9.4%	NA	NA
Sun Microsystems	JAVA	David C. Bailey	\$	16.38	-23.3%	-10.7%	-27.8%	-33.8%
Symantec Corp.	SYMC	Sarah Friar	\$	17.09	-13.4%	2.7%	-11.5%	-2.6%
Synchronoss Technologies Inc.	SNCR	Elizabeth W. Gausam, CFA	\$	20.49	-40.3%	-45.9%	-52.5%	11.5%
Syniverse Technologies, Inc.	SVR	Elizabeth W. Gausam, CFA	\$	17.40	32.1%	7.4%	8.2%	59.6%
Taleo Corporation	TLEO	Sasa Zorovic, Ph.D.	\$	19.55	-9.2%	-41.1%	-21.9%	16.2%
Tech Data Corp.	TECD	Min Park	\$	33.04	-15.6%	-13.1%	-16.7%	-8.6%
Tellabs, Inc.	TLAB	Brantley Thompson	\$	5.99	-49.5%	-13.3%	-35.9%	-41.4%
Teradyne, Inc.	TER	James Covello	\$	12.53	-29.6%	19.1%	-10.4%	-26.4%
Texas Instruments Inc.	TXN	James Covello	\$	29.29	-25.0%	-14.2%	-19.2%	-6.6%
TIBCO Software Inc.	TIBX	Derek Bingham	\$	7.69	-16.0%	-5.8%	3.4%	-16.0%
Total System Services, Inc.	TSS	Elizabeth W. Gausam, CFA	\$	23.32	-21.1%	-14.0%	-15.3%	-27.8%
Trident Microsystems, Inc.	TRID	James Schneider, Ph.D.	\$	5.54	-67.0%	-9.2%	-62.9%	-74.1%
Unisys Corporation	UIS	Julio C. Quinteros, Jr.	\$	4.70	-48.9%	-7.8%	-30.4%	-44.3%
VeriFone Holdings, Inc.	PAY	Elizabeth W. Gausam, CFA	\$	16.94	-54.1%	-26.4%	-58.8%	-55.2%
Vengy Ltd.	VRGY	James Covello	\$	18.10	-36.0%	-31.7%	-31.9%	-24.7%
VenSign, Inc.	VRSN	Sarah Friar	\$	35.60	7.9%	-6.2%	4.5%	40.5%
Volterra Semiconductor Corp.	VLTR	Simona Jankowski, CFA	\$	10.83	-19.8%	-8.6%	-14.1%	-18.9%
Western Digital Corp.	WDC	David C. Bailey	\$	30.55	44.0%	0.9%	20.6%	73.3%
Western Union Co.	WU	Elizabeth W. Gausam, CFA	\$	21.80	5.6%	-10.4%	3.5%	-2.4%
Wipro Ltd. (ADR)	WIT	Julio C. Quinteros, Jr.	\$	11.56	-25.8%	-24.8%	-16.8%	-30.9%
WNS (Holdings) Ltd.	WNS	Julio C. Quinteros, Jr.	\$	16.03	-43.1%	-1.2%	-4.1%	-45.7%
XILINX Corp.	XLNX	James Schneider, Ph.D.	\$	24.94	-10.9%	11.1%	-6.5%	-8.0%
Zoran Corporation	ZRAN	James Schneider, Ph.D.	\$	14.15	-30.6%	-39.0%	-30.5%	-18.7%
S&P 500				1352.99	-12.7%	-9.6%	-10.8%	-5.8%

Note: This table shows movement in absolute share price and not total shareholder return.
Results presented should not and cannot be viewed as an indicator of future performance.

Source: FactSet, Quantum database.

March 26, 2008

Americas: Communications Technology

Financial Advisory disclosure

Goldman Sachs is acting as financial advisor to Motorola, Inc. in an announced strategic transaction.

March 26, 2008

Americas: Communications Technology

Reg AC

We, Simona Jankowski, CFA and Thomas D. Lee, hereby certify that all of the views expressed in this report accurately reflect our personal views about the subject company or companies and its or their securities. We also certify that no part of our compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Investment profile

The Goldman Sachs Investment Profile provides investment context for a security by comparing key attributes of that security to its peer group and market. The four key attributes depicted are: growth, returns, multiple and volatility. Growth, returns and multiple are indexed based on composites of several methodologies to determine the stocks percentile ranking within the region's coverage universe.

The precise calculation of each metric may vary depending on the fiscal year, industry and region but the standard approach is as follows:

Growth is a composite of next year's estimate over current year's estimate, e.g. EPS, EBITDA, Revenue. **Return** is a year one prospective aggregate of various return on capital measures, e.g. CROCI, ROACE, and ROE. **Multiple** is a composite of one-year forward valuation ratios, e.g. P/E, dividend yield, EV/FCF, EV/EBITDA, EV/DACF, Price/Book. **Volatility** is measured as trailing twelve-month volatility adjusted for dividends.

Quantum

Quantum is Goldman Sachs' proprietary database providing access to detailed financial statement histories, forecasts and ratios. It can be used for in-depth analysis of a single company, or to make comparisons between companies in different sectors and markets.

Disclosures

Coverage group(s) of stocks by primary analyst(s)

Simona Jankowski, CFA: America-Semi Device, America-TelecomEqpt. Thomas D. Lee: America-TelecomEqpt.

America-Semi Device: Advanced Micro Devices, Inc., Altera Corp., Analog Devices, Inc., Applied Micro Circuits, Broadcom Corporation, Intel Corp., International Rectifier Corp., Intersil Corp., Lattice Semiconductor Corp., Leadis Technology, Linear Technology Corp., LSI Corp., Marvell Technology Group Ltd., Maxim Integrated Products, Micrel, Inc., Microchip Technology Inc., Micron Technology Inc., Monolithic Power Systems, National Semiconductor Corp., Nvidia Corp., PMC-Sierra, Inc., SanDisk Corporation, Sigma Designs, Inc., Silicon Image, Inc., SiRF Technology Holdings, Inc., Texas Instruments Inc., Trident Microsystems, Inc., Volterra Semiconductor Corp., XILINX Corp., Zoran Corporation.

America-TelecomEqpt: Acme Packet, Inc., Aruba Networks, Inc., Cisco Systems, Inc., Juniper Networks, Inc., Motorola, Inc., Netgear, Inc., Powerwave Technologies, Inc., QUALCOMM, Inc., Research In Motion Ltd., Research In Motion Ltd., Starent Networks, Corp..

Company-specific regulatory disclosures

The following disclosures relate to relationships between The Goldman Sachs Group, Inc. (with its affiliates, "Goldman Sachs") and companies covered by the Global Investment Research Division of Goldman Sachs and referred to in this research.

Goldman Sachs beneficially owned 1% or more of common equity (excluding positions managed by affiliates and business units not required to be aggregated under US securities law) as of the month end preceding this report: Netgear, Inc. (\$20.66)

Goldman Sachs has received compensation for investment banking services in the past 12 months: Acme Packet, Inc. (\$7.85), Aruba Networks, Inc. (\$6.07), Cisco Systems, Inc. (\$24.78), Motorola, Inc. (\$10.02) and Starent Networks, Corp. (\$14.22)

Goldman Sachs expects to receive or intends to seek compensation for investment banking services in the next 3 months: Acme Packet, Inc. (\$7.85), Aruba Networks, Inc. (\$6.07), Cisco Systems, Inc. (\$24.78), Juniper Networks, Inc. (\$26.20), Motorola, Inc. (\$10.02), Netgear, Inc. (\$20.66), Powerwave Technologies, Inc. (\$2.74), QUALCOMM, Inc. (\$40.49), Research In Motion Ltd. (\$118.15) and Starent Networks, Corp. (\$14.22)

Goldman Sachs has received compensation for non-investment banking services during the past 12 months: Acme Packet, Inc. (\$7.85), Cisco Systems, Inc. (\$24.78), Juniper Networks, Inc. (\$26.20), Motorola, Inc. (\$10.02), QUALCOMM, Inc. (\$40.49) and Research In Motion Ltd. (\$118.15)

Goldman Sachs had an investment banking services client relationship during the past 12 months with: Acme Packet, Inc. (\$7.85), Aruba Networks, Inc. (\$6.07), Cisco Systems, Inc. (\$24.78), Juniper Networks, Inc. (\$26.20), Motorola, Inc. (\$10.02), QUALCOMM, Inc. (\$40.49), Research In Motion Ltd. (\$118.15) and Starent Networks, Corp. (\$14.22)

Goldman Sachs had a non-investment banking securities-related services client relationship during the past 12 months with: Acme Packet, Inc. (\$7.85), Aruba Networks, Inc. (\$6.07), Cisco Systems, Inc. (\$24.78), Juniper Networks, Inc. (\$26.20), Motorola, Inc. (\$10.02), QUALCOMM, Inc. (\$40.49) and Research In Motion Ltd. (\$118.15)

Goldman Sachs had a non-securities services client relationship during the past 12 months with: Cisco Systems, Inc. (\$24.78), Juniper Networks, Inc. (\$26.20), Motorola, Inc. (\$10.02), QUALCOMM, Inc. (\$40.49) and Research In Motion Ltd. (\$118.15)

Goldman Sachs has managed or co-managed a public or Rule 144A offering in the past 12 months: Motorola, Inc. (\$10.02) and Starent Networks, Corp. (\$14.22)

March 26, 2008

Americas: Communications Technology

A director and/or employee of Goldman Sachs is a director: QUALCOMM, Inc. (\$40.49)

Goldman Sachs makes a market in the securities or derivatives thereof: Acme Packet, Inc. (\$7.85), Aruba Networks, Inc. (\$6.07), Cisco Systems, Inc. (\$24.78), Juniper Networks, Inc. (\$26.20), Motorola, Inc. (\$10.02), Netgear, Inc. (\$20.66), Powerwave Technologies, Inc. (\$2.74), QUALCOMM, Inc. (\$40.49), Research In Motion Ltd. (\$118.15) and Starent Networks, Corp. (\$14.22)

Goldman Sachs is a specialist in the relevant securities and will at any given time have an inventory position, "long" or "short," and may be on the opposite side of orders executed on the relevant exchange: Acme Packet, Inc. (\$7.85), Aruba Networks, Inc. (\$6.07), Cisco Systems, Inc. (\$24.78) and Netgear, Inc. (\$20.66)

Distribution of ratings/investment banking relationships

Goldman Sachs Investment Research global coverage universe

	Rating Distribution			Investment Banking Relationships		
	Buy	Hold	Sell	Buy	Hold	Sell
Global	29%	58%	13%	39%	30%	26%

As of Jan 1, 2008, Goldman Sachs Global Investment Research had investment ratings on 2,994 equity securities. Goldman Sachs assigns stocks as Buys and Sells on various regional Investment Lists; stocks not so assigned are deemed Neutral. Such assignments equate to Buy, Hold and Sell for the purposes of the above disclosure required by NASD/NYSE rules. See 'Ratings, Coverage groups and views and related definitions' below.

Price target and rating history chart(s)

Compendium report: please see disclosures at <http://www.gs.com/research/hedge.html>. Disclosures applicable to the companies included in this compendium can be found in the latest relevant published research.

Regulatory disclosures

Disclosures required by United States laws and regulations

See company-specific regulatory disclosures above for any of the following disclosures required as to companies referred to in this report: manager or co-manager in a pending transaction; 1% or other ownership; compensation for certain services; types of client relationships; managed/co-managed public offerings in prior periods; directorships; market making and/or specialist role.

The following are additional required disclosures: **Ownership and material conflicts of interest:** Goldman Sachs policy prohibits its analysts, professionals reporting to analysts and members of their households from owning securities of any company in the analyst's area of coverage. **Analyst compensation:** Analysts are paid in part based on the profitability of Goldman Sachs, which includes investment banking revenues. **Analyst as officer or director:** Goldman Sachs policy prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director, advisory board member or employee of any company in the analyst's area of coverage. **Distribution of ratings:** See the distribution of ratings disclosure above. **Price chart:** See the price chart, with changes of ratings and price targets in prior periods, above, or, if electronic format or if with respect to multiple companies which are the subject of this report, on the Goldman Sachs website at <http://www.gs.com/research/hedge.html>. Goldman, Sachs & Co. is a member of SIPC.

Additional disclosures required under the laws and regulations of jurisdictions other than the United States

The following disclosures are those required by the jurisdiction indicated, except to the extent already made above pursuant to United States laws and regulations. **Australia:** This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act. **Canada:** Goldman Sachs Canada Inc. has approved of, and agreed to take responsibility for, this research in Canada if and to the extent it relates to equity securities of Canadian issuers. Analysts may conduct site visits but are prohibited from accepting payment or reimbursement by the company of travel expenses for such visits. **Hong Kong:** Further information on the securities of covered companies referred to in this research may be obtained on request from Goldman Sachs (Asia) L.L.C. **India:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (India) Securities Private Limited; **Japan:** See below. **Korea:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (Asia) L.L.C., Seoul Branch. **Russia:** Research reports distributed in the Russian Federation are not advertising as defined in Russian law, but are information and analysis not having product promotion as their main purpose and do not provide appraisal within the meaning of the Russian Law on Appraisal. **Singapore:** Further information on the covered companies referred to in this research may be obtained from Goldman Sachs (Singapore) Pte. (Company Number: 198602165W). **United Kingdom:** Persons who would be categorized as retail clients in the United Kingdom, as such term is defined in the rules of the Financial Services Authority, should read this research in conjunction with prior Goldman Sachs research on the covered companies referred to herein and should refer to the risk warnings that have been sent to them by Goldman Sachs International. A copy of these risks warnings, and a glossary of certain financial terms used in this report, are available from Goldman Sachs International on request.

European Union: Disclosure information in relation to Article 4 (1) (d) and Article 6 (2) of the European Commission Directive 2003/126/EC is available at http://www.gs.com/client_services/global_investment_research/europeanpolicy.html

Japan: Goldman Sachs Japan Co., Ltd. is a Financial Instrument Dealer under the Financial Instrument and Exchange Law, registered with the Kanto Financial Bureau (Registration No. 69), and is a member of Japan Securities Dealers Association (JSDA) and Financial Futures Association of Japan (FFJA). Sales and purchase of equities are subject to commission pre-determined with clients plus consumption tax. See company-specific disclosures as to any applicable disclosures required by Japanese stock exchanges, the Japanese Securities Dealers Association or the Japanese Securities Finance Company.

Ratings, coverage groups and views and related definitions

Buy (B), Neutral (N), Sell (S) -Analysts recommend stocks as Buys or Sells for inclusion on various regional Investment Lists. Being assigned a Buy or Sell on an Investment List is determined by a stock's return potential relative to its coverage group as described below. Any stock not assigned as a Buy or a Sell on an Investment List is deemed Neutral. Each regional Investment Review Committee manages various regional Investment Lists to

March 26, 2008

Americas: Communications Technology

a global guideline of 25%-35% of stocks as Buy and 10%-15% of stocks as Sell; however, the distribution of Buys and Sells in any particular coverage group may vary as determined by the regional Investment Review Committee. Regional Conviction Buy and Sell lists represent investment recommendations focused on either the size of the potential return or the likelihood of the realization of the return.

Return potential represents the price differential between the current share price and the price target expected during the time horizon associated with the price target. Price targets are required for all covered stocks. The return potential, price target and associated time horizon are stated in each report adding or reiterating an Investment List membership.

Coverage groups and views: A list of all stocks in each coverage group is available by primary analyst, stock and coverage group at <http://www.gs.com/research/hedge.html>. The analyst assigns one of the following coverage views which represents the analyst's investment outlook on the coverage group relative to the group's historical fundamentals and/or valuation. **Attractive (A).** The investment outlook over the following 12 months is favorable relative to the coverage group's historical fundamentals and/or valuation. **Neutral (N).** The investment outlook over the following 12 months is neutral relative to the coverage group's historical fundamentals and/or valuation. **Cautious (C).** The investment outlook over the following 12 months is unfavorable relative to the coverage group's historical fundamentals and/or valuation.

Not Rated (NR). The investment rating and target price, if any, have been removed pursuant to Goldman Sachs policy when Goldman Sachs is acting in an advisory capacity in a merger or strategic transaction involving this company and in certain other circumstances. **Rating Suspended (RS).** Goldman Sachs Research has suspended the investment rating and price target, if any, for this stock, because there is not a sufficient fundamental basis for determining an investment rating or target. The previous investment rating and price target, if any, are no longer in effect for this stock and should not be relied upon. **Coverage Suspended (CS).** Goldman Sachs has suspended coverage of this company. **Not Covered (NC).** Goldman Sachs does not cover this company. **Not Available or Not Applicable (NA).** The information is not available for display or is not applicable. **Not Meaningful (NM).** The information is not meaningful and is therefore excluded.

Ratings, coverage views and related definitions prior to June 26, 2006

Our rating system requires that analysts rank order the stocks in their coverage groups and assign one of three investment ratings (see definitions below) within a ratings distribution guideline of no more than 25% of the stocks should be rated Outperform and no fewer than 10% rated Underperform. The analyst assigns one of three coverage views (see definitions below), which represents the analyst's investment outlook on the coverage group relative to the group's historical fundamentals and valuation. Each coverage group, listing all stocks covered in that group, is available by primary analyst, stock and coverage group at <http://www.gs.com/research/hedge.html>.

Definitions

Outperform (OP). We expect this stock to outperform the median total return for the analyst's coverage universe over the next 12 months. **In-Line (IL).** We expect this stock to perform in line with the median total return for the analyst's coverage universe over the next 12 months. **Underperform (U).** We expect this stock to underperform the median total return for the analyst's coverage universe over the next 12 months.

Coverage views: **Attractive (A).** The investment outlook over the following 12 months is favorable relative to the coverage group's historical fundamentals and/or valuation. **Neutral (N).** The investment outlook over the following 12 months is neutral relative to the coverage group's historical fundamentals and/or valuation. **Cautious (C).** The investment outlook over the following 12 months is unfavorable relative to the coverage group's historical fundamentals and/or valuation.

Current Investment List (CIL). We expect stocks on this list to provide an absolute total return of approximately 15%-20% over the next 12 months. We only assign this designation to stocks rated Outperform. We require a 12-month price target for stocks with this designation. Each stock on the CIL will **automatically** come off the list after 90 days unless renewed by the covering analyst and the relevant Regional Investment Review Committee.

Global product; distributing entities

The Global Investment Research Division of Goldman Sachs produces and distributes research products for clients of Goldman Sachs, and pursuant to certain contractual arrangements, on a global basis. Analysts based in Goldman Sachs offices around the world produce equity research on industries and companies, and research on macroeconomics, currencies, commodities and portfolio strategy.

This research is disseminated in Australia by Goldman Sachs JBWere Pty Ltd (ABN 21 006 797 897) on behalf of Goldman Sachs; in Canada by Goldman Sachs Canada Inc. regarding Canadian equities and by Goldman Sachs & Co. (all other research); in Germany by Goldman Sachs & Co. oHG; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs JBWere (NZ) Limited on behalf of Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman, Sachs & Co. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom and European Union.

European Union: Goldman Sachs International, authorised and regulated by the Financial Services Authority, has approved this research in connection with its distribution in the European Union and United Kingdom; Goldman, Sachs & Co. oHG, regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht, may also be distributing research in Germany.

General disclosures in addition to specific disclosures required by certain jurisdictions

This research is for our clients only. Other than disclosures relating to Goldman Sachs, this research is based on current public information that we consider reliable, but we do not represent it is accurate or complete, and it should not be relied on as such. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Other than certain industry reports published on a periodic basis, the large majority of reports are published at irregular intervals as appropriate in the analyst's judgment.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management, and brokerage business. We have investment banking and other business relationships with a substantial percentage of the companies covered by our Global Investment Research Division.

Our salespeople, traders, and other professionals may provide oral or written market commentary or trading strategies to our clients and our proprietary trading desks that reflect opinions that are contrary to the opinions expressed in this research. Our asset management area, our proprietary trading desks and investing businesses may make investment decisions that are inconsistent with the recommendations or views expressed in this research.

March 26, 2008

Americas: Communications Technology

We and our affiliates, officers, directors, and employees, excluding equity analysts, will from time to time have long or short positions in, act as principal in, and buy or sell, the securities or derivatives (including options and warrants) thereof of covered companies referred to in this research.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Clients should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The price and value of the investments referred to in this research and the income from them may fluctuate. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. Certain transactions, including those involving futures, options, and other derivatives, give rise to substantial risk and are not suitable for all investors.

Current options disclosure documents are available from Goldman Sachs sales representatives or at <http://www.theocc.com/publications/risks/riskchap1.jsp>. Fluctuations in exchange rates could have adverse effects on the value or price of, or income derived from, certain investments.

Our research is disseminated primarily electronically, and, in some cases, in printed form. Electronic research is simultaneously available to all clients.

Disclosure information is also available at <http://www.gs.com/research/hedge.html> or from Research Compliance, One New York Plaza, New York, NY 10004.

Copyright 2008 The Goldman Sachs Group, Inc.

No part of this material may be (i) copied, photocopied or duplicated in any form by any means or (ii) redistributed without the prior written consent of The Goldman Sachs Group, Inc.

EXHIBIT 26

North America Equity Research
07 April 2008



Aruba Networks

Downgrading to Underweight as Economic Downturn Continues to Take its Toll; Reducing Ests

We are downgrading our rating on Aruba to Underweight from Neutral as we believe extending sales cycles and shrinking order sizes could continue to plague Aruba throughout 08 – especially in the retailing vertical – since the economic downturn appears longer than we previously expected. And as Aruba invests to expand internationally and to develop its 2-tier distribution system through partner training, we believe it could take longer to achieve its 19-20% LT op margin target. So despite Aruba having emerged as a #2 WLAN player to Cisco, we rate the stock Underweight.

- **We expect elongated sales cycles and shrinking order sizes to plague Aruba through much of 08**, matching comments from Riverbed last week, and resulting in an increased probability of future earnings misses.
- **Lowering our 08 growth forecast for the dependent WLAN market 5 points to 28%**, down from our previous estimate of 33% as we believe WLAN deployments are among the more discretionary IT projects, implying that 802.11n roll-outs could be delayed until 2009.
- **On the flip side, Aruba appears to be gaining more share** – an add'l point in each of the last two qtrs – however, a new crop of WLAN startups could make it difficult for ARUN to gain share going forward.
- **We are lowering our FQ308 revenue and EPS estimates to \$41.7M and (\$0.02)** from \$42.6M and (\$0.01), the low end of guidance.
- **Lowering FY08 and FY09 revenue ests to \$174M and \$220M** from \$177M and \$235M on our lower market forecast, as our EPS estimates fall to \$0.05 and \$0.16 from \$0.07 and \$0.23, all ex-stock comp.
- **Aruba trades at 23.9x our CY09 EPS estimate of \$0.22** (ex-stock comp), a premium to its peer group average of 12.4x.

Aruba Networks, Inc. (ARUN;ARUN US)

	2007A	2008E (Old)	2008E (New)	2009E (Old)	2009E (New)
EPS (\$)					
Q1 (Oct)	(0.28)	(0.02)A	(0.02)A	0.00	(0.02)
Q2 (Jan)	(0.35)	(0.05)A	(0.05)A	0.01	0.01
Q3 (Apr)	(0.09)	(0.06)	(0.07)	0.02	(0.01)
Q4 (Jul)	(0.04)	(0.02)	(0.04)	0.03	0.00
FY	(0.42)	(0.14)	(0.17)	0.05	(0.02)
Revenues FY (\$ mn)	127	177	174	235	220

Note: Estimates in earnings box are post-stock compensation expense. EPS estimates pre-stock comp are \$0.05 in FY08 and \$0.16 in FY09 and \$0.08 in CY08 and \$0.22 in CY09.
Source: Company data, Reuters, JPMorgan estimates.

Downgrade Underweight

Previous Rating: Neutral

\$5.26

04 April 2008

Communications Equipment & Data Networking

Ehud A. Gelblum, PhD^{AC}

(1-212) 622-6457

ehud.gelblum@jpmorgan.com

Kimberly A. Watkins, CFA

(1-415) 315-6724

kimberly.a.watkins@jpmorgan.com

Price Performance



Source: RIMES, Reuters.

Company Data	
Price (\$)	5.26
Date Of Price	04 Apr 08
52-week Range (\$)	23.85 - 4.63
Mkt Cap (\$ mn)	479.71
Fiscal Year End	Jul
Shares O/S (mn)	91

www.morganmarkets.com

J.P. Morgan Securities Inc.

See page 5 for analyst certification and important disclosures, including investment banking relationships.

JPMorgan does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Customers of JPMorgan in the United States can receive independent, third-party research on the company or companies covered in this report, at no cost to them, where such research is available. Customers can access this independent research at www.morganmarkets.com or can call 1-800-477-0406 toll free to request a copy of this research.

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008



Investment Thesis & Conclusion

We are downgrading Aruba to Underweight from Neutral, as we believe extending sales cycles and shrinking order sizes could slow down the entire WLAN market and thereby continue to plague Aruba throughout '08 - especially in the retailing vertical, which is one of its four largest along with Federal government, healthcare, and education – since we believe the economic downturn appears poised to last well into 2009. In addition, as Aruba continues to invest in expanding internationally and in developing its 2-tier distribution system through partner training, we believe it could take longer to achieve its 19-20% LT op margin target. So despite Aruba having emerged as a solid #2 WLAN player to Cisco, we rate the stock Underweight.

Lowering our WLAN Market Forecasts

Lowering our '08 growth forecast for the dependent WLAN market by 5 points to 28%, down from our previous 33% estimate, as we believe WLAN deployments are among the most discretionary of IT projects. Therefore, we now believe that 802.11n roll-outs, which the market is counting on to re-accelerate growth, could be delayed until 2009. This revision takes our market size estimates to \$1,389M in CY08, down from \$1,527M. We are also lowering our CY09 market growth to 24% from 29.6%, taking our market size estimates to \$1,722M in CY09, down from \$1,979M.

Table 1: Lowering our Dependent WLAN Market Growth Forecast in '08 to 28%

\$ in millions

	Fiscal 2008E	Fiscal 2009E	Calendar 2008E	Calendar 2009E
Total Dependent WLAN Market - NEW	1,207.7	1,548.6	1,388.6	1,722.1
y/y Market Growth	26.6%	28.2%	27.7%	24.0%
Aruba Brand Mkt Share %	10.6%	10.9%	10.3%	11.1%
Alcatel Brand Mkt Share %	1.6%	1.3%	1.2%	1.2%
Total Aruba Mkt Share %	12.2%	12.2%	11.6%	12.3%
Aruba Brand Dependent Revenue	128.2	168.0	143.4	190.6
Alcatel Brand Dependent Revenue	19.0	20.1	17.3	21.3
Total Aruba Dependent Revenue	147.2	188.2	160.7	211.9
Total Dependent WLAN Market - OLD	1,321.0	1,745.7	1,526.6	1,978.8
y/y Market Growth	38.8%	32.1%	33.1%	29.6%
Aruba Brand Mkt Share %	9.9%	10.4%	9.9%	10.3%
Alcatel Brand Mkt Share %	1.3%	1.3%	1.2%	1.2%
Total Aruba Mkt Share %	11.2%	11.7%	11.1%	11.5%
Aruba Brand Dependent Revenue	131.4	180.7	151.7	203.5
Alcatel Brand Dependent Revenue	16.9	22.8	18.4	24.2
Total Aruba Dependent Revenue	148.3	203.5	170.1	227.8
Total Dependent WLAN Market - CHANGE	(113.3)	(197.1)	(138.1)	(256.7)
y/y Market Growth (bps)	(1230 bps)	(390 bps)	(540 bps)	(560 bps)
Aruba Brand Mkt Share (bps)	70 bps	50 bps	40 bps	80 bps
Alcatel Brand Mkt Share %	30 bps	---	---	---
Total Aruba Mkt Share %	100 bps	50 bps	40 bps	80 bps
Aruba Brand Dependent Revenue	(3.2)	(12.6)	(8.3)	(12.9)
Alcatel Brand Dependent Revenue	2.1	(2.7)	(1.1)	(2.9)
Total Aruba Dependent Revenue	(1.1)	(15.3)	(9.3)	(15.9)

Source: Company reports, Dell'Oro Group, and JPMorgan estimates.

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008



Estimate Revisions

We are lowering our FQ308 revenue and EPS to \$41.7M and (\$0.02) from \$42.6M and (\$0.01), the low end of guidance. We are also lowering FY08 and FY09 revenue to \$174M and \$220M from \$177M and \$235M, respectively, on the slower market growth, while EPS in the same years falls to \$0.05 and \$0.16 from \$0.07 and \$0.23, all ex-stock comp. See Table 2 below for more detail on our estimate revisions.

Table 2: Aruba Estimate Revisions

\$ in millions

	Apr-08 Q3-08E	Jul-08 Q4-08E	Oct-08 Q1-09E	Jan-09 Q2-09E	Apr-09 Q3-09E	Jul-09 Q4-09E	2008E	2009E
New revenue	41.7	45.2	48.2	55.7	55.4	60.6	174.3	220.0
Old revenue	<u>42.9</u>	<u>47.1</u>	<u>52.0</u>	<u>58.3</u>	<u>60.9</u>	<u>64.2</u>	<u>177.3</u>	<u>235.4</u>
\$M change	(1.1)	(1.9)	(3.7)	(2.6)	(5.5)	(3.6)	(3.1)	(15.4)
New Gross Margin	67.9%	67.0%	67.0%	67.0%	67.1%	67.1%	68.0%	67.1%
Old Gross Margin	<u>67.9%</u>	<u>67.0%</u>	<u>67.0%</u>	<u>67.0%</u>	<u>67.1%</u>	<u>67.1%</u>	<u>68.0%</u>	<u>67.0%</u>
% Change	---	---	---	---	---	---	---	---
New Operating Margin	-4.6%	0.6%	3.5%	7.7%	5.5%	6.3%	0.4%	5.9%
Old Operating Margin	<u>-2.7%</u>	<u>3.3%</u>	<u>6.0%</u>	<u>7.7%</u>	<u>8.6%</u>	<u>9.9%</u>	<u>1.6%</u>	<u>8.2%</u>
% Change	(200 bps)	(270 bps)	(250 bps)	---	(310 bps)	(350 bps)	(120 bps)	(230 bps)
New EPS (ex. stock comp)	(\$0.02)	\$0.01	\$0.03	\$0.05	\$0.04	\$0.04	\$0.05	\$0.16
Old EPS (ex. stock comp)	<u>(\$0.01)</u>	<u>\$0.03</u>	<u>\$0.04</u>	<u>\$0.05</u>	<u>\$0.06</u>	<u>\$0.07</u>	<u>\$0.07</u>	<u>\$0.23</u>
\$ change	(\$0.01)	(\$0.01)	(\$0.02)	(\$0.00)	(\$0.02)	(\$0.03)	(\$0.02)	(\$0.07)
New EPS (post stock comp)	(\$0.07)	(\$0.04)	(\$0.02)	\$0.01	(\$0.01)	\$0.00	(\$0.17)	(\$0.02)
Old EPS (post stock comp)	<u>(\$0.06)</u>	<u>(\$0.02)</u>	<u>(\$0.00)</u>	<u>\$0.01</u>	<u>\$0.02</u>	<u>\$0.03</u>	<u>(\$0.14)</u>	<u>\$0.05</u>
\$ change	(\$0.01)	(\$0.02)	(\$0.02)	(\$0.00)	(\$0.02)	(\$0.03)	(\$0.03)	(\$0.07)

Source: Company reports and JPMorgan estimates.

Valuation and Rating Analysis

Aruba Networks has emerged as the leading #2 player and primary alternative to Cisco in the enterprise WLAN equipment market, a market that we believe is still in its early innings, but also one of the most discretionary in an IT budget, and therefore subject to trimming during economic downturns such as the one we are currently experiencing. So although Aruba continues to gain share, we believe elongated sales cycles and smaller order sizes could result in earnings misses, driving our Underweight rating with the stock trading at a rich 23.9x our CY09 EPS estimate of \$0.22 (ex-stock compensation expense), a premium to its peer group at 12.4x.

Risks to Rating

Potential upside risks to our Underweight rating include the potential the WLAN market grows faster than we expect, that the market is immune to the economic downturn, that competition is more mild than we anticipate and that Aruba is acquired. An additional upside risk is that the relationship with OEM partner Alcatel-Lucent produces greater growth for Aruba than we expect.

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008



Table 3: Aruba Networks Income Statement

\$ in millions

	Oct-06 Q1-07A	Jan-07 Q2-07A	Apr-07 Q3-07A	Jul-07 Q4-07A	Oct-07 Q1-08A	Jan-08 Q2-08A	Apr-08 Q3-08E	Jul-08 Q4-08E	Oct-08 Q1-09E	Jan-09 Q2-09E	Apr-09 Q3-09E	Jul-09 Q4-09E	2006A	2007A	2008E	2009E
Product Revenue	19,106	22,662	29,777	36,394	38,458	34,170	34,417	37,761	40,850	47,806	47,465	52,123	43,170	107,939	144,806	188,244
Service Revenue	2,121	2,657	3,816	4,254	7,273	5,548	6,500	6,700	7,800	7,600	7,752	8,372	2,985	12,847	26,021	30,724
Ratable Revenue	3,278	1,329	1,068	1,038	0,999	0,927	0,800	0,700	0,400	0,300	0,200	0,100	26,347	6,713	3,426	1,000
Total Revenue	24,505	26,647	34,661	41,686	46,730	40,645	41,717	45,161	48,250	55,706	55,417	60,595	72,503	127,469	174,253	219,968
Total COGS	9,614	10,105	10,394	12,928	14,886	12,547	13,389	14,898	15,903	18,357	18,257	19,958	29,851	43,041	55,720	72,476
Product Gross Profit	11,827	14,148	20,878	25,183	26,701	23,244	23,404	25,300	27,369	32,030	31,802	34,922	26,272	72,036	98,648	126,124
Service Gross Profit	0,972	1,570	2,718	2,919	4,506	4,257	4,420	4,523	4,725	5,130	5,233	5,651	0,599	8,179	17,706	20,739
Ratable Gross Profit	2,092	0,824	0,671	0,656	0,637	0,597	0,504	0,441	0,252	0,189	0,126	0,063	15,781	4,243	2,179	0,630
Total Gross Profit	14,891	16,542	24,267	28,758	31,844	28,098	28,328	30,263	32,346	37,349	37,160	40,636	42,651	84,458	118,533	147,492
Product Gross Margin	61.9%	62.4%	70.1%	69.2%	69.4%	68.0%	68.0%	67.0%	67.0%	67.0%	67.0%	67.0%	60.9%	66.7%	68.1%	67.0%
Service Gross Margin	45.8%	59.1%	71.2%	68.6%	62.0%	76.7%	68.0%	67.5%	67.5%	67.5%	67.5%	67.5%	20.1%	63.7%	68.0%	67.5%
Ratable Gross Margin	63.8%	62.0%	62.8%	63.2%	63.8%	64.4%	63.0%	63.0%	63.0%	63.0%	63.0%	63.0%	528.6%	33.0%	8.4%	2.1%
Total Gross Margin	60.8%	62.1%	70.0%	69.0%	68.1%	69.1%	67.9%	67.0%	67.0%	67.0%	67.1%	67.1%	58.8%	66.2%	68.0%	67.1%
R & D Expense	4,865	5,279	6,032	6,552	6,850	7,768	8,145	7,298	7,720	7,966	8,285	9,007	13,871	22,728	30,061	32,978
Selling & Marketing	9,919	11,883	15,199	18,753	19,005	17,148	18,862	19,305	19,541	21,168	21,588	23,160	33,015	55,754	74,320	85,458
General and Administrative	1,960	2,554	2,390	2,592	3,281	3,495	3,258	3,390	3,377	3,899	4,259	4,632	5,750	9,496	13,424	16,168
Other	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000
Total Operating Expenses	16,744	19,716	23,621	27,897	29,136	28,411	30,265	29,993	30,639	33,034	34,133	36,799	52,636	87,979	117,805	134,604
Operating Income	(1,854)	(3,174)	0,646	0,861	2,708	(0,313)	(1,937)	0,270	1,708	4,315	3,027	3,837	(9,985)	(3,521)	0,728	12,888
Operating Margin	-7.6%	-11.9%	1.9%	2.1%	5.8%	-0.8%	-4.6%	0.6%	3.5%	7.7%	5.5%	6.3%	-13.8%	-2.8%	0.4%	5.9%
Depreciation & Amortization	0,396	0,446	0,542	0,624	0,903	0,975	1,310	1,403	1,506	1,615	1,766	1,883	1,549	2,008	4,591	6,770
EBITDA	(1,458)	(2,728)	1,188	1,485	3,611	0,662	(0,627)	1,672	3,214	5,930	4,793	5,721	(8,436)	(1,513)	5,318	19,658
EBITDA Margin	-5.9%	-10.2%	3.4%	3.6%	7.7%	1.6%	-1.5%	3.7%	6.7%	10.6%	8.6%	9.4%	-11.6%	-1.2%	3.1%	8.9%
Interest Income	0,112	0,197	0,567	1,344	1,356	1,264	1,110	1,131	1,124	1,122	1,124	1,141	0,551	2,219	4,860	4,511
Interest (Expense)	(0,035)	(0,028)	(0,022)	(0,003)	0,000	0,000	0,000	0,000	0,000	0,000	0,000	0,000	(0,315)	(0,088)	0,000	0,000
Other Income / (Expense)	(0,047)	(0,059)	(0,061)	(0,085)	0,011	(0,144)	(0,100)	(0,100)	(0,100)	(0,100)	(0,200)	(0,200)	(0,097)	(0,252)	(0,333)	(0,600)
Earnings Before Taxes	(1,824)	(3,064)	1,130	2,117	4,075	0,807	(0,928)	1,300	2,732	5,337	3,951	4,779	(9,647)	(1,641)	5,255	16,799
Income Tax Provision / (Benefit)	0,088	0,123	0,082	0,081	0,224	0,228	0,250	0,300	0,400	0,450	0,500	0,550	0,306	0,374	1,002	1,900
Effective Tax Rate	-4.8%	-4.0%	7.3%	3.8%	5.5%	28.3%	-27.0%	23.1%	14.6%	8.4%	12.7%	11.5%	-3.1%	-22.8%	19.1%	11.3%
Net Income to Common	(1,912)	(3,187)	1,048	2,036	3,851	0,579	(1,178)	1,000	2,332	4,887	3,451	4,229	(10,153)	(2,015)	4,253	14,899
Stock Option Comp Expense	(1,814)	(1,689)	(4,460)	(4,754)	(5,205)	(4,093)	(4,093)	(4,093)	(4,093)	(4,093)	(4,093)	(4,093)	(1,255)	(12,717)	(17,484)	(16,372)
Net Income Post Stock Comp	(3,726)	(4,876)	(3,412)	(2,718)	(1,354)	(3,514)	(5,271)	(3,093)	(1,761)	0,794	(0,642)	0,136	(11,408)	(14,732)	(13,231)	(1,473)
EPS - Diluted	(\$0.14)	(\$0.23)	\$0.01	\$0.02	\$0.04	\$0.01	(\$0.02)	\$0.01	\$0.03	\$0.05	\$0.04	\$0.04	(\$0.91)	(\$0.06)	\$0.05	\$0.16
EPS - Diluted, Post Stock Comp	(\$0.28)	(\$0.35)	(\$0.09)	(\$0.04)	(\$0.02)	(\$0.05)	(\$0.07)	(\$0.04)	(\$0.02)	\$0.01	(\$0.01)	\$0.00	(\$1.02)	(\$0.42)	(\$0.17)	(\$0.02)
Share Outstanding - Diluted (M)	63,538	70,101	82,297	93,496	94,167	91,254	90,000	91,125	92,264	93,417	94,585	95,767	64,993	77,358	91,637	94,008

Sources: Company reports and JPMorgan estimates.

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008




Analyst Certification:

The research analyst(s) denoted by an "AC" on the cover of this report certifies (or, where multiple research analysts are primarily responsible for this report, the research analyst denoted by an "AC" on the cover or within the document individually certifies, with respect to each security or issuer that the research analyst covers in this research) that: (1) all of the views expressed in this report accurately reflect his or her personal views about any and all of the subject securities or issuers; and (2) no part of any of the research analyst's compensation was, is, or will be directly or indirectly related to the specific recommendations or views expressed by the research analyst(s) in this report.

Important Disclosures

- **Market Maker:** JPMSI makes a market in the stock of Aruba Networks.
- **Client of the Firm:** Aruba Networks is or was in the past 12 months a client of JPMSI; during the past 12 months, JPMSI provided to the company investment banking services.
- **Investment Banking (past 12 months):** JPMSI or its affiliates received in the past 12 months compensation for investment banking services from Aruba Networks.
- **Investment Banking (next 3 months):** JPMSI or its affiliates expect to receive, or intend to seek, compensation for investment banking services in the next three months from Aruba Networks.

Aruba Networks (ARUN) Price Chart



Date	Rating	Share Price (\$)	Price Target (\$)
13-Aug-07	N	18.95	-

Source: Reuters and JPMorgan; price data adjusted for stock splits and dividends.
Initiated coverage Aug 13, 2007. This chart shows JPMorgan's continuing coverage of this stock; the current analyst may or may not have covered it over the entire period.
JPMorgan ratings: OW = Overweight, N = Neutral, UW = Underweight.

Explanation of Equity Research Ratings and Analyst(s) Coverage Universe:

JPMorgan uses the following rating system: **Overweight** [Over the next six to twelve months, we expect this stock will outperform the average total return of the stocks in the analyst's (or the analyst's team's) coverage universe.] **Neutral** [Over the next six to twelve months, we expect this stock will perform in line with the average total return of the stocks in the analyst's (or the analyst's team's) coverage universe.] **Underweight** [Over the next six to twelve months, we expect this stock will underperform the average total return of the stocks in the analyst's (or the analyst's team's) coverage universe.] The analyst or analyst's team's coverage universe is the sector and/or country shown on the cover of each publication. See below for the specific stocks in the certifying analyst(s) coverage universe.

Coverage Universe: **Ehud A. Gelblum, PhD:** ADTRAN, Inc (ADTN), ARRIS Group (ARRS), Acme Packet (APKT), Aruba Networks (ARUN), CIENA Corp (CIEN), Cisco Systems (CSCO), F5 Networks (FFIV), Infinera Corporation (INFN), JDSU (JDSU), Juniper Networks (JNPR), Motorola (MOT), Nortel Networks (NT), Opnext (OPXT), QUALCOMM (QCOM), ShoreTel (SHOR), Starent Networks (STAR), Tellabs (TLAB)

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008



JPMorgan Equity Research Ratings Distribution, as of March 31, 2008

	Overweight (buy)	Neutral (hold)	Underweight (sell)
JPM Global Equity Research Coverage	46%	41%	13%
IB clients*	50%	51%	39%
JPMI Equity Research Coverage	41%	47%	11%
IB clients*	70%	66%	54%

*Percentage of investment banking clients in each rating category.

For purposes only of NASD/NYSE ratings distribution rules, our Overweight rating falls into a buy rating category; our Neutral rating falls into a hold rating category; and our Underweight rating falls into a sell rating category.

Valuation and Risks: Please see the most recent company-specific research report for an analysis of valuation methodology and risks on any securities recommended herein. Research is available at <http://www.morganmarkets.com>, or you can contact the analyst named on the front of this note or your JPMorgan representative.

Analysts' Compensation: The equity research analysts responsible for the preparation of this report receive compensation based upon various factors, including the quality and accuracy of research, client feedback, competitive factors, and overall firm revenues, which include revenues from, among other business units, Institutional Equities and Investment Banking.

Other Disclosures

Options related research: If the information contained herein regards options related research, such information is available only to persons who have received the proper option risk disclosure documents. For a copy of the Option Clearing Corporation's Characteristics and Risks of Standardized Options, please contact your JPMorgan Representative or visit the OCC's website at <http://www.optionsclearing.com/publications/risks/riskstoc.pdf>.

Legal Entities Disclosures

U.S.: JPMSI is a member of NYSE, FINRA and SIPC. J.P. Morgan Futures Inc. is a member of the NFA. JPMorgan Chase Bank, N.A. is a member of FDIC and is authorized and regulated in the UK by the Financial Services Authority. **U.K.:** J.P. Morgan Securities Ltd. (JPMSL) is a member of the London Stock Exchange and is authorised and regulated by the Financial Services Authority. Registered in England & Wales No. 2711006. Registered Office 125 London Wall, London EC2Y 5AJ. **South Africa:** J.P. Morgan Equities Limited is a member of the Johannesburg Securities Exchange and is regulated by the FSB. **Hong Kong:** J.P. Morgan Securities (Asia Pacific) Limited (CE number AAJ321) is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission in Hong Kong. **Korea:** J.P. Morgan Securities (Far East) Ltd, Seoul branch, is regulated by the Korea Financial Supervisory Service. **Australia:** J.P. Morgan Australia Limited (ABN 52 002 888 011/AFS Licence No: 238188) is regulated by ASIC and J.P. Morgan Securities Australia Limited (ABN 61 003 245 234/AFS Licence No: 238066) is a Market Participant with the ASX and regulated by ASIC. **Taiwan:** J.P. Morgan Securities (Taiwan) Limited is a participant of the Taiwan Stock Exchange (company-type) and regulated by the Taiwan Securities and Futures Bureau. **India:** J.P. Morgan India Private Limited is a member of the National Stock Exchange of India Limited and The Stock Exchange, Mumbai and is regulated by the Securities and Exchange Board of India. **Thailand:** JPMorgan Securities (Thailand) Limited is a member of the Stock Exchange of Thailand and is regulated by the Ministry of Finance and the Securities and Exchange Commission. **Indonesia:** PT J.P. Morgan Securities Indonesia is a member of the Jakarta Stock Exchange and Surabaya Stock Exchange and is regulated by the BAPEPAM. **Philippines:** J.P. Morgan Securities Philippines Inc. is a member of the Philippine Stock Exchange and is regulated by the Securities and Exchange Commission. **Brazil:** Banco J.P. Morgan S.A. is regulated by the Comissão de Valores Mobiliários (CVM) and by the Central Bank of Brazil. **Mexico:** J.P. Morgan Casa de Bolsa, S.A. de C.V., J.P. Morgan Grupo Financiero is a member of the Mexican Stock Exchange and authorized to act as a broker dealer by the National Banking and Securities Exchange Commission. **Singapore:** This material is issued and distributed in Singapore by J.P. Morgan Securities Singapore Private Limited (JPMSS) [mica (p) 207/01/2008 and Co. Reg. No.: 199405335R] which is a member of the Singapore Exchange Securities Trading Limited and is regulated by the Monetary Authority of Singapore (MAS) and/or JPMorgan Chase Bank, N.A., Singapore branch (JPMCB Singapore) which is regulated by the MAS. **Malaysia:** This material is issued and distributed in Malaysia by JPMorgan Securities (Malaysia) Sdn Bhd (18146-x) which is a Participating Organization of Bursa Malaysia Securities Bhd and is licensed as a dealer by the Securities Commission in Malaysia. **Pakistan:** J. P. Morgan Pakistan Broking (Pvt.) Ltd is a member of the Karachi Stock Exchange and regulated by the Securities and Exchange Commission of Pakistan.

Country and Region Specific Disclosures

U.K. and European Economic Area (EEA): Issued and approved for distribution in the U.K. and the EEA by JPMSL. Investment research issued by JPMSL has been prepared in accordance with JPMSL's Policies for Managing Conflicts of Interest in Connection with Investment Research which outline the effective organisational and administrative arrangements set up within JPMSL for the prevention and avoidance of conflicts of interest with respect to research recommendations, including information barriers, and can be found at <http://www.jpmorgan.com/pdfdoc/research/ConflictManagementPolicy.pdf>. This report has been issued in the U.K. only to persons of a kind described in Article 19 (5), 38, 47 and 49 of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (all such persons being referred to as "relevant persons"). This document must not be acted on or relied on by persons who are not relevant persons. Any investment or investment activity to which this document relates is only available to relevant persons and will be engaged in only with relevant persons. In other EEA countries, the report has been issued to persons regarded as professional investors (or equivalent) in their home jurisdiction. **Germany:**

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008



This material is distributed in Germany by J.P. Morgan Securities Ltd. Frankfurt Branch and JPMorgan Chase Bank, N.A., Frankfurt Branch who are regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht. **Australia:** This material is issued and distributed by JPMSAL in Australia to "wholesale clients" only. JPMSAL does not issue or distribute this material to "retail clients." The recipient of this material must not distribute it to any third party or outside Australia without the prior written consent of JPMSAL. For the purposes of this paragraph the terms "wholesale client" and "retail client" have the meanings given to them in section 761G of the Corporations Act 2001. **Hong Kong:** The 1% ownership disclosure as of the previous month end satisfies the requirements under Paragraph 16.5(a) of the Hong Kong Code of Conduct for persons licensed by or registered with the Securities and Futures Commission. (For research published within the first ten days of the month, the disclosure may be based on the month end data from two months' prior.) J.P. Morgan Broking (Hong Kong) Limited is the liquidity provider for derivative warrants issued by J.P. Morgan International Derivatives Ltd and listed on The Stock Exchange of Hong Kong Limited. An updated list can be found on HKEx website: <http://www.hkex.com.hk/prod/dw/Lp.htm>. **Japan:** There is a risk that a loss may occur due to a change in the price of the shares in the case of share trading, and that a loss may occur due to the exchange rate in the case of foreign share trading. In the case of share trading, JPMorgan Securities Japan Co., Ltd., will be receiving a brokerage fee and consumption tax (shouhizei) calculated by multiplying the executed price by the commission rate which was individually agreed between JPMorgan Securities Japan Co., Ltd., and the customer in advance. Financial Instruments Firms: JPMorgan Securities Japan Co., Ltd., Kanto Local Finance Bureau (kinsho) No. [82] Participating Association / Japan Securities Dealers Association, The Financial Futures Association of Japan. **Korea:** This report may have been edited or contributed to from time to time by affiliates of J.P. Morgan Securities (Far East) Ltd, Seoul branch. **Singapore:** JPMSI and/or its affiliates may have a holding in any of the securities discussed in this report; for securities where the holding is 1% or greater, the specific holding is disclosed in the Legal Disclosures section above. **India:** For private circulation only not for sale. **Pakistan:** For private circulation only not for sale. **New Zealand:** This material is issued and distributed by JPMSAL in New Zealand only to persons whose principal business is the investment of money or who, in the course of and for the purposes of their business, habitually invest money. JPMSAL does not issue or distribute this material to members of "the public" as determined in accordance with section 3 of the Securities Act 1978. The recipient of this material must not distribute it to any third party or outside New Zealand without the prior written consent of JPMSAL.

General: Additional information is available upon request. Information has been obtained from sources believed to be reliable but JPMorgan Chase & Co. or its affiliates and/or subsidiaries (collectively JPMorgan) do not warrant its completeness or accuracy except with respect to any disclosures relative to JPMSI and/or its affiliates and the analyst's involvement with the issuer that is the subject of the research. All pricing is as of the close of market for the securities discussed, unless otherwise stated. Opinions and estimates constitute our judgment as of the date of this material and are subject to change without notice. Past performance is not indicative of future results. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The opinions and recommendations herein do not take into account individual client circumstances, objectives, or needs and are not intended as recommendations of particular securities, financial instruments or strategies to particular clients. The recipient of this report must make its own independent decisions regarding any securities or financial instruments mentioned herein. JPMSI distributes in the U.S. research published by non-U.S. affiliates and accepts responsibility for its contents. Periodic updates may be provided on companies/industries based on company specific developments or announcements, market conditions or any other publicly available information. Clients should contact analysts and execute transactions through a JPMorgan subsidiary or affiliate in their home jurisdiction unless governing law permits otherwise.

"Other Disclosures" last revised February 6, 2008.

Copyright 2008 JPMorgan Chase & Co. All rights reserved. This report or any portion hereof may not be reprinted, sold or redistributed without the written consent of JPMorgan.

Ehud A. Gelblum, PhD
(1-212) 622-6457
ehud.gelblum@jpmorgan.com

North America Equity Research
07 April 2008

JPMorgan 



EXHIBIT 27

**JAYWALK CONSENSUS 4/21/08 3.42**

Independent Research Providers Covering
Aruba Networks, Incorporated ARUN

ABOUT THE JAYWALK CONSENSUS

The Jaywalk Consensus is an average of all of the independent research providers' ratings on the given security. By averaging these ratings, investors are given insight into the independent research community's perspective on individual securities. The research providers who participate in the Jaywalk Consensus are professional firms that attest to having no investment banking or other potential conflicts that might impact the integrity of their research.

CURRENT JAYWALK CONSENSUS on 4/21/08**ARUN STOCK PRICE HISTORY****KEY STATISTICS**

as of 04/18/2008

Industry	Computer Peripherals
Sector	TECHNOLOGY
Close Price	\$5.22
52-Week Price Range	\$4.40 - \$23.85
Volume	1,035,300
Average Daily Volume	1,293,300
Market Capitalization	\$420.31 Million
Enterprise Value	\$354.90 Million
P/E (Trailing)	--
Dividend	--
Dividend Yield	0.0%
5-year Historical EPS Growth Rate	--
Liquidity Ratio	\$1,439
Shares Outstanding	80.5 Million
Employees	441


BREAKDOWN OF INDEPENDENT RATINGS

Below is the breakdown of analyst ratings in the Jaywalk Consensus. For example, if there is a Jaywalk Consensus of 2.5 with 10 IRPs covering the stock, and there are 5 buys and 5 holds, this shows a solid overall consensus with no negative ratings. Additionally, if this stock had a consensus of 3.0 last month and a consensus of 4.0 two months ago, then the consensus on that stock shows a trend of consistent improvement.

	4/21/08	3/31/08	2/29/08	1/31/08	12/31/07
Very Positive or Strong Buy (Rating=1)	1	0	0	0	0
Positive or Buy (Rating=2)	0	0	0	2	2
Neutral or Hold (Rating=3)	6	7	4	3	3
Negative or Sell (Rating=4)	3	4	3	2	2
Very Negative or Strong Sell (Rating=5)	2	2	2	2	2
Total Independent Research Providers	12	13	9	9	9
Jaywalk Consensus	3.42	3.62	3.78	3.44	3.44

INDEPENDENT RESEARCH PROVIDER RATING DISTRIBUTION

12 Independent Research Providers


**INDUSTRY, SECTOR AND OVERALL JAYWALK UNIVERSE RANKING**

Industry, Sector and Overall Jaywalk Universe ranking charts on the right are shown on a 0.0% to 100.0% scale. A stock with a **percent rank in industry** of 25%, for example, indicates approximately 75% of the stocks in that industry have a better Jaywalk Consensus than the given stock. Likewise, a **percent rank in sector** of 50% indicates that the stock is ranked the midpoint of all stocks within the sector. Rollover dots in the chart to view company and consensus information.

INDUSTRY CONSENSUS DISTRIBUTION

This chart shows ARUN's Consensus rank in industry relative to the Consensus of the 35 companies in its industry. ARUN is shown against its peers to assist investors in identifying where the Jaywalk Consensus and coverage of the stock is in respect to other industry participants.

Note: Dots displayed indicate there are at least 4 IRPs providing coverage.

**PEER COMPARISON****RANK IN INDUSTRY 36.1%**

Computer Peripherals
 ARUN Coverage: 12 IRPs, Industry Avg.: 11 IRPs

0% 100%
 ARUN's consensus score is lower than 63.9% of the Industry.

RANK IN SECTOR 23.0%

TECHNOLOGY
 ARUN Coverage: 12 IRPs, Sector Avg.: 12 IRPs

0% 100%
 ARUN's consensus score is lower than 77.0% of the Sector.

RANK IN JAYWALK UNIVERSE 20.4%

ARUN Coverage: 12 IRPs, Universe Avg.: 11 IRPs

0% 100%
 ARUN's consensus score is lower than 79.6% of the Jaywalk Universe.

Important Information The information contained herein is for informational purposes only and is not intended to provide tax, legal or investment advice. This information is compiled directly from (1) the research providers' recommendations on the equity securities displayed and (2) CoreData, a third-party financial data vendor and is based on information sources believed to be accurate and reliable. Such information, however, is presented without warranty of any kind. Past performance should not be taken as an indication of future performance, and no representation of any kind is made as to future performance. BNY Jaywalk and CoreData are not soliciting any action based upon such information, or endorsing any recommendation or opinion expressed by the Jaywalk Consensus and the research providers. Accordingly, neither BNY Jaywalk and any of its affiliates, nor CoreData, guarantee the timeliness, accuracy, completeness, sequence, or adequacy of the information. You agree that any and all use you make of the information is solely at your own risk and without recourse to BNY Jaywalk, any of its affiliates, CoreData, or any independent research provider.

EXHIBIT 28



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
 United States Patent and Trademark Office
 P.O. Box 1460
 Alexandria, VA 22313-1460
 www.uspto.gov

Inter Partes Reexamination Filing Data - June 30, 2007

1. Total requests filed since start of *inter partes* reexam on 11/29/99 272¹
2. Number of filings by discipline
 - a. Chemical Operation 69 25%
 - b. Electrical Operation 98 36%
 - c. Mechanical Operation 105 39%
3. Annual Reexam Filings

<u>Fiscal Yr.</u>	<u>No.</u>	<u>Fiscal Yr.</u>	<u>No.</u>	<u>Fiscal Yr.</u>	<u>No.</u>	<u>Fiscal Yr.</u>	<u>No.</u>
2000	0	2002	4	2004	27	2006	70
2001	1	2003	21	2005	59	2007	90
4. Number known to be in litigation.....137.....50%
5. Decisions on requests 232
 - a. No. granted 223 96%
 - (1) By examiner 2238
 - (2) By Director (on petition) 0
 - b. No. not granted 9 4%
 - (1) By examiner 7
 - (2) Reexam vacated 2
6. Overall reexamination pendency (Filing date to certificate issue date)
 - a. Average pendency 29.1 (mos.)
 - b. Median pendency 30.4 (mos.)
7. Total *inter partes* reexamination certificates issued (1999 - present) 8
 - a. Certificates with all claims confirmed 1 12%
 - b. Certificates with all claims canceled 7 88%
 - c. Certificates with claims changes 0 0%

¹Of the requests received through June 30, 2007, 1 proceeding was vacated per 37 CFR 1.913; and 7 requests have not yet been accorded a filing date and preprocessing of 1 request was terminated, for failure to comply with the requirements of 37 CFR 1.915. See Clarification of Filing Date Requirements for *Ex Parte* and *Inter Partes* Reexamination Proceedings, Final Rule, 71 Fed. Reg. 44219 (August 4, 2006).



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
 United States Patent and Trademark Office
 P.O. Box 1450
 Alexandria, VA 22313-1450
www.uspto.gov

Ex Parte Reexamination Filing Data - December 31, 20071. Total requests filed since start of ex parte reexam on 07/01/81 9060¹

a. By patent owner	3495	39%
b. By other member of public	5400	59%
c. By order of Commissioner	165	2%

2. Number of filings by discipline

a. Chemical Operation	2703	30%
b. Electrical Operation	3023	33%
c. Mechanical Operation	3334	37%

3. Annual Ex Parte Reexam Filings

Fiscal Yr.	No.	Fiscal Yr.	No.	Fiscal Yr.	No.	Fiscal Yr.	No.
1981	78 (3 mos.)	1989	243	1997	376	2005	524
1982	187	1990	297	1998	350	2006	511
1983	186	1991	307	1999	385	2007	643
1984	189	1992	392	2000	318	2008	165
1985	230	1993	359	2001	296		
1986	232	1994	379	2002	272		
1987	240	1995	392	2003	392		
1988	268	1996	418	2004	441		

4. Number known to be in litigation 2398 26%

5. Determinations on requests 8714

a. No. granted 7998 92%

(1) By examiner	7885
(2) By Director (on petition)	113

b. No. denied 716 8%

(1) By examiner	681
(2) Order vacated	35

¹Of the requests received in FY 2008, 23 requests have not yet been accorded a filing date, and preprocessing of 3 requests was terminated for failure to comply with the requirements of 37 CFR 1.510. See Clarification of Filing Date Requirements for *Ex Parte* and *Inter Partes* Reexamination Proceedings, Final Rule, 71 Fed. Reg. 44219 (August 4, 2006).

6. Total examiner denials (includes denials reversed by Director)	794			
a. Patent owner requester	439		55%	
b. Third party requester	355		45%	
7. Overall reexamination pendency (Filing date to certificate issue date)				
a. Average pendency			24.0 (mos.)	
b. Median pendency			18.6 (mos.)	
8. Reexam certificate claim analysis:	Owner Requester	3rd Party Requester	Comm'r Initiated	Overall
a. All claims confirmed	23%	29%	12%	26%
b. All claims cancelled	7%	12%	21%	10%
c. Claims changes	70%	59%	67%	64%
9. Total ex parte reexamination certificates issued (1981 - present)	6066			
a. Certificates with all claims confirmed	1556	26%		
b. Certificates with all claims canceled	636	10%		
c. Certificates with claims changes	3874	64%		
10. Reexam claim analysis - requester is patent owner or 3rd party; or Comm'r initiated.				
a. Certificates - PATENT OWNER REQUESTER	2607			
(1) All claims confirmed	592	23%		
(2) All claims canceled	194	7%		
(3) Claim changes	1821	70%		
b. Certificates - 3rd PARTY REQUESTER	3313			
(1) All claims confirmed	946	29%		
(2) All claims canceled	413	12%		
(3) Claim changes	1954	59%		
c. Certificates - COMM'R INITIATED REEXAM	146			
(1) All claims confirmed	18	12%		
(2) All claims canceled	30	21%		
(3) Claim changes	98	67%		

EXHIBIT 29



US006973622B1

(12) **United States Patent**
Rappaport et al.

(10) **Patent No.:** **US 6,973,622 B1**(45) **Date of Patent:** **Dec. 6, 2005**

(54) **SYSTEM AND METHOD FOR DESIGN, TRACKING, MEASUREMENT, PREDICTION AND OPTIMIZATION OF DATA COMMUNICATION NETWORKS**

(75) Inventors: **Theodore Rappaport**, Salem, VA (US);
Roger Skidmore, Blacksburg, VA (US);
Benjamin Henty, Blacksburg, VA (US)

(73) Assignee: **Wireless Valley Communications, Inc.**,
Austin, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 724 days.

(21) Appl. No.: **09/668,145**

(22) Filed: **Sep. 25, 2000**

(51) Int. Cl.⁷ **G06F 3/00**; G06F 19/00;
G06F 15/16

(52) U.S. Cl. **715/735**; 715/736; 703/21;
703/22; 709/221

(58) Field of Search 703/2, 3, 5, 21,
703/22; 455/33.1, 33.4, 564, 446; 345/133;
202/186; 715/735, 736, 734; 709/221, 222

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,675,147 A 6/1987 Schaefer et al.
4,736,453 A 4/1988 Schloemer
4,885,694 A 12/1989 Pray et al.
5,111,392 A 5/1992 Malin
5,119,307 A 6/1992 Blaha et al.
5,239,487 A 8/1993 Horejsi et al.
5,293,640 A 3/1994 Gunmar et al.
5,307,261 A 4/1994 Maki et al.
5,337,149 A 8/1994 Kozah et al.
5,339,184 A 8/1994 Tang
5,375,123 A 12/1994 Andersson et al.
5,394,522 A 2/1995 Sanchez-Frank et al.
5,450,615 A 9/1995 Fortune et al.
5,458,123 A 10/1995 Unger
5,465,390 A 11/1995 Cohen

5,467,441 A 11/1995 Stone et al.
5,482,050 A 1/1996 Smokoff et al.
5,485,568 A 1/1996 Venable et al.
5,491,644 A 2/1996 Pickering et al.
5,491,837 A 2/1996 Haartsen
5,493,679 A 2/1996 Virgil et al.
5,515,269 A 5/1996 Willis et al.
5,528,518 A 6/1996 Bradshaw et al.
5,539,665 A 7/1996 Lamming et al.
5,553,312 A 9/1996 Gattey et al.
5,553,620 A 9/1996 Snider et al.
5,555,354 A 9/1996 Strasnick et al.
5,561,841 A 10/1996 Markus
5,564,070 A 10/1996 Want et al.
5,586,254 A 12/1996 Kondo
5,594,946 A 1/1997 Menich et al.
5,598,532 A 1/1997 Liron
5,625,827 A 4/1997 Krause et al.
5,636,344 A 6/1997 Lewis
5,689,355 A 11/1997 Okubo et al.

(Continued)

OTHER PUBLICATIONS

Article "Building Database Manipulator" Copyright, Jan. 1998; MPRG and Virginia Tech.

(Continued)

Primary Examiner—Larry D. Donaghue

(74) *Attorney, Agent, or Firm*—Whitham, Curtis & Christofferson, PC


(57)

ABSTRACT

A system and method for design, tracking, measurement, prediction and optimization of data communications networks includes a site specific model of the physical environment, and performs a wide variety of different calculations for predicting network performance using a combination of prediction modes and measurement data based on the components used in the communications networks, the physical environment, and radio propagation characteristics.

68 Claims, 6 Drawing Sheets

Method for optimizing a data communications network using predictions and measurements



US 6,973,622 B1

Page 2

U.S. PATENT DOCUMENTS

5,710,758 A 1/1998 Soliman et al.
 5,755,072 A 5/1998 Lingafelter
 5,761,093 A 6/1998 Urbish et al.
 5,774,669 A 6/1998 George et al.
 5,794,128 A 8/1998 Brockel et al.
 5,799,154 A 8/1998 Kuriyan
 5,802,146 A 9/1998 Dulman
 5,809,282 A 9/1998 Cooper et al.
 5,815,395 A 9/1998 Hart et al.
 5,821,937 A 10/1998 Tonelli et al.
 5,825,759 A 10/1998 Liu
 5,828,960 A 10/1998 Tang et al.
 5,831,610 A 11/1998 Tonelli et al.
 5,832,389 A 11/1998 Dent
 5,845,124 A 12/1998 Berman
 5,861,887 A 1/1999 Butler et al.
 5,867,112 A 2/1999 Kost
 5,877,777 A 3/1999 Colwell
 5,878,328 A 3/1999 Chawla et al.
 5,907,850 A 5/1999 Krause et al.
 5,917,808 A 6/1999 Kosbab
 5,923,850 A 7/1999 Barroux
 5,926,762 A 7/1999 Arpee et al.
 5,940,196 A 8/1999 Piehler et al.
 5,945,976 A 8/1999 Iwamura et al.
 5,948,055 A 9/1999 Pulsipher et al.
 5,949,335 A 9/1999 Maynard
 5,949,988 A 9/1999 Feisullin et al.
 5,953,669 A 9/1999 Stratis et al.
 5,963,867 A 10/1999 Reynolds et al.
 5,970,406 A 10/1999 Komara
 5,977,851 A 11/1999 Stancil et al.
 5,987,328 A 11/1999 Ephremides et al.
 5,994,984 A 11/1999 Stancil et al.
 6,006,021 A 12/1999 Tognazzini
 6,018,625 A 1/2000 Hayball et al.
 6,021,316 A 2/2000 Heiska et al.
 6,032,105 A 2/2000 Lee et al.
 6,038,547 A 3/2000 Casto
 6,044,273 A 3/2000 Tekinay
 6,058,102 A 5/2000 Drysdale et al.
 6,058,262 A 5/2000 Kawas et al.
 6,059,842 A 5/2000 Dumarot et al.
 6,061,722 A 5/2000 Lipa et al.
 6,075,541 A 6/2000 Maclinovsky
 6,085,335 A 7/2000 Djoko et al.
 6,088,522 A 7/2000 Lee et al.
 6,104,699 A 8/2000 Holender et al.
 6,108,309 A 8/2000 Cohoe et al.
 6,111,857 A 8/2000 Soliman et al.
 6,122,083 A 9/2000 Ohta et al.
 6,148,010 A 11/2000 Sutton et al.
 6,199,032 B1 3/2001 Anderson
 6,204,813 B1 3/2001 Wadell et al.
 6,208,833 B1 3/2001 Preschutti et al.
 6,229,540 B1 5/2001 Tonelli et al.
 6,243,772 B1 6/2001 Ghori et al.
 6,253,086 B1 6/2001 Parantainen et al.
 6,285,377 B1 9/2001 Greenbaum et al.
 6,289,203 B1 9/2001 Smith et al.
 6,311,144 B1 10/2001 Abu El Ata
 6,317,599 B1* 11/2001 Rappaport et al. 455/446
 6,326,987 B2 12/2001 Alexander
 6,330,005 B1 12/2001 Tonelli et al.
 6,337,688 B1 1/2002 Berstis
 6,338,031 B1 1/2002 Lee et al.
 6,356,758 B1 3/2002 Almeida et al.
 6,393,432 B1 5/2002 Flansburg et al.
 6,408,312 B1 6/2002 Forthman et al.
 6,442,507 B1* 8/2002 Skidmore et al. 702/186
 6,470,195 B1 10/2002 Meyer

6,487,417 B1 11/2002 Rossoni et al.
 6,493,679 B1* 12/2002 Rappaport et al. 705/29
 6,496,290 B1 12/2002 Lee
 6,499,006 B1* 12/2002 Rappaport et al. 703/20
 6,505,045 B1 1/2003 Hills et al.
 6,625,454 B1* 9/2003 Rappaport et al. 455/446

OTHER PUBLICATIONS

PCS 97 Track 7; Engineering & Systems Management; T. Rappaport.
 Propagator; vol. 8, No. 3; Fall 1997.
 SMT Plus 1.0 User's Manual; R. Skidmore & T. Rappaport; Copyright, Aug. 1996; Virginia Tech.
 Software by Andrew titled "RF Planner" dated Jun. 17, 1997.
 A user guide titled "Andrew Microwave System Planner" dated Jul. 1999.
 A user guide titled "Andrew Antenna System Panner" dated Jun. 1999.
 From Bird's Eye Real-time Mapping Software dated Jun. 30, 2002.
 IEEE Transactions on Antennas and propagation, vol. 46, No. 8, Aug. 1998. "Effect of Terrain on Path Loss in Urban Enviroments for Wireless Applications" Leonard Piazzi and Henry L. Bertoni.
 P. Bahl, V. Padmanabhan, and A. Balachandran, "A Software System for Locating Mobile Users: Design, Evaluation, and Lessons," Microsoft Technical Report, Apr. 2000.
 G. Durgin, T.S. Rappaport, H. Xu, Measurements and Models for Radio Path Loss and Penetration Loss in and Around Homes and Trees at 5.85 GHz, IEEE Transactions on Communications, vol. 46, No. 11, Nov. 1998.
 C.M. Peter Ho et al., "Antenna Effects on Indoor Obstructed Wireless Channels and a Deterministic Image-Based Wide-Band Propagation Model for In-Building Personal Communications Systems," International Journal of Wireless Information Networks, vol. 1, No. 1, 1994.
 S. Kim et al., "Radio Propagation Measurements and Predictions Using Three-dimensional Ray Tracing in Urban Enviroments at 908 MHz and 1.9 GHz," IEEE Transactions on Vehicular Technology, vol. 48, No. 3, May 1999.
 T.S., Rappaport et al., "Use of Topographic Maps with Building Information to Determine Antenna Placements and GPS Satellite Coverage for Radio Detection and Tracking in Urban Environments," MPRG Technical Reports MPRG-TR-95-14, Virginia Tech, Sep. 1995.
 R.K. Morrow, Jr. and T.S. Rappaport, "Getting In," Wireless Review Magazine, Mar. 2000.
 Wireless Valley Communications, Inc., "SitePlanner 3.16 for Windows 95/98/NT User's Manual," Software User's Manual, pp. 5-148 to 5-156, 1999.
 M. Panjwani et al., "Interactive Computation of Coverage Regions for Wireless Communication in Multifloored Indoor Environments," IEEE Journal on Selected Areas in Communications, vol. 14, No. 3, Apr. 1996.
 L. Piazzi and H.L. Bertoni, "Achievable Accuracy of Site-Specific Path-Loss Predictions in Residential Enviroments" IEEE Transactions on Vehicular Technology, vol. 48, No. 3, May 1999.
 T.S. Rappaport et al., "Wireless Communication: Past Events and a Future Perspective", IEEE Communications Magazine, May 2002.
 T.S. Rappaport et al., "Radio Propagation Prediction Techniques and Computer-Aided Channeling Modeling for Embedded Wireless Microsystems," ARPA Annual Report,

US 6,973,622 B1

Page 3

- MPRG Technical Report MPRG-TR-94-12, Virginia Tech, Jul. 1994.
- T.S., Rappaport et al., "Use of Topographic Maps with Building Information to Determine Antenna Placements for Radio Detection and Tracking in Urban Environments," MPRG Technical Report MPRG-TR-95-14, Virginia Tech, Nov. 1995.
- D. Ullmo et al., "Wireless Propagation in Buildings: A Statistical Scattering Approach," IEEE Transactions on Vehicular Technology, vol. 48, No. 3, May 1999.
- T.S. Rappaport, "wireless Communications: Principles and Practice" Second Edition, Prentice Hall, 2002.
- T.S., Rappaport et al., "Use of Topographic Maps with Building Information to Determine Antenna Placements and GPS Satellite Coverage for Radio Detection and Tracking in Urban Environments," MPRG Technical Report MPRG-TR-95-14, Virginia Tech, Sep. 1995.
- T.S. Rappaport et al., "Indoor Path Loss Measurement for Homes and Apartments at 2.4 and 5.85 GHz," private report produced for Motorola, Dec. 16, 1997.
- T.S. Rappaport, "Isolating Interference," Wireless Review Magazine, May 2000.
- Slides from T.S. Rappaport and R. Skidmore, "Introduction to In-Building Wireless Systems," Infocast In-Building Wireless Solutions Conference and Exposition, Feb. 4, 2003.
- S. Sandhu, M.P. Koushik, and T.S. Rappaport "Predicted Path Loss for Roslyn VA, First set of predictions for ORD Projection Site Specific Propagation Prediction," MPRG Technical Report MPRG-TR-94-20, Virginia Tech, Dec. 1994.
- S. Sandhu, M.P. Koushik, and T.S. Rappaport "Predicted Path Loss for Roslyn VA, First set of predictions for ORD Projection Site Specific Propagation Prediction," MPRG Technical Report MPRG-TR-94-20, Virginia Tech, Mar. 1995.
- S. Seidel et al., "Site-Specific Propagation Prediction for Wireless In-Building Personal Communications Design," IEEE Transactions on Vehicular Technology, vol. 43, No. 4, Nov. 1994.
- S. Shakkottai and T.S. Rappaport, "Research Challenges in Wireless Networks: A Technical Overview," Proceeding of the Fifth International Symposium on Wireless Personal Multimedia Communications, Honolulu, HI, Oct. 2002.
- H. Sherali et al., "On the Optimal Location of Transmitters for Micro-cellular Radio Communication System Design," IEEE Journal on Selected Areas in Communications, vol. 14, No. 3, pp. 662-673, May. 1996.
- R. Skidmore et al., "A Comprehensive In-Building and Microcellular Wireless Communication System Design Tool" The Bradley Department of Electrical Engineering, MPRG-TR-97-13, Jun. 1997. Master's Thesis—unpublished by Virginia Tech for 2 years after submission.
- R. Skidmore, et al., Russell Senate Office Building Propagation Study, Project Report for Joseph R. Loring & Associates; "Project Update," AoC Contract #Aabr96088, prepared for Office of the Architect of Capital, Jan. 19, 1997.
- R. Skidmore, et al., Russell Senate Office Building Propagation Study, Project Report for Joseph R. Loring & Associates; "Assessment and Study of the Proposed Enhancements of the Wireless Communications Environment of the Russell Senate Office Building (RSOB) and Associated Utility Tunnels," AoC Contract #Aabr96088, prepared for Office of the Architect of the Capitol, Feb. 20, 1997.
- R. Torres et al., "CINDOOR: An Engineering Tool for Planning and Design of Wireless Systems in Enclosed Spaces," IEEE Antennas and Propagation Magazine, vol. 41, No. 4 Aug. 1999.
- R. Skidmore et al., "Interactive Coverage Region and System Design Simulation for Wireless Communication Systems in Multi-Floored Indoor Environments: SMT Plus tm," IEEE ICUPC Proceedings, 1996.
- T.S. Rappaport et al., "Radio Propagation Prediction Techniques and Computer-Aided Channel Modeling for Embedded Wireless Microsystems," MPRG Tech. Report MPRG-TR-95-08, Virginia Tech, Jul. 1995.
- Company Web Page "Actix" www.actix.com product name: E-NOS (now E-AMS), no date given.
- Company Web Page "Agilent" www.agilent.com product name: OPAS32, no date given.
- Company Web Page "Agilent" www.agilent.com product name: Wizard, no date given.
- Company Web Page "Camarco" www.edx.com product name: SignalPro, no date given.
- Company Web Page "ComOpt" www.comopt.com product name: Cellopt AFP, no date given.
- Company Web Page "Lucent" www.bell-labs.com product name: WISE, no date given.
- Company Web Page "Ericsson" www.ericsson.com product name: TEMS Lite, no date given.
- Company Web Page "Ericsson" www.ericsson.com product name: TEMS, no date given.
- Company Web Page "Marconi" www.marconi.com product name: PlanET, no date given.
- Company Web Page "Marconi" www.marconi.com product name: deciblePlanner, no date given.
- Company Web Page "Schema" www.schema.com product name: Optimizer, no date given.
- Company Web Page "ScoreBoard" www.scoreboard.com product name: ScoreBoard, no date given.

* cited by examiner

U.S. Patent

Dec. 6, 2005

Sheet 1 of 6

US 6,973,622 B1

Figure 1: Example transmission of data over a communications network





Figure 2: Creation of a digital signal from an analog signal



U.S. Patent

Dec. 6, 2005

Sheet 2 of 6

US 6,973,622 B1

Figure 3: Illustration of the difference between bits, packets and frames.


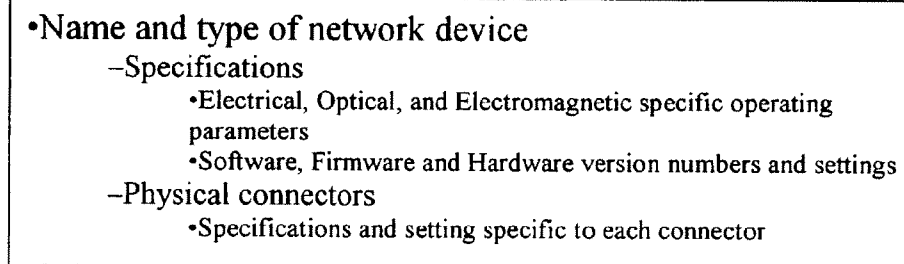


Figure 4: Illustration of the data displayed in each node of the Tree View of a data communications network.



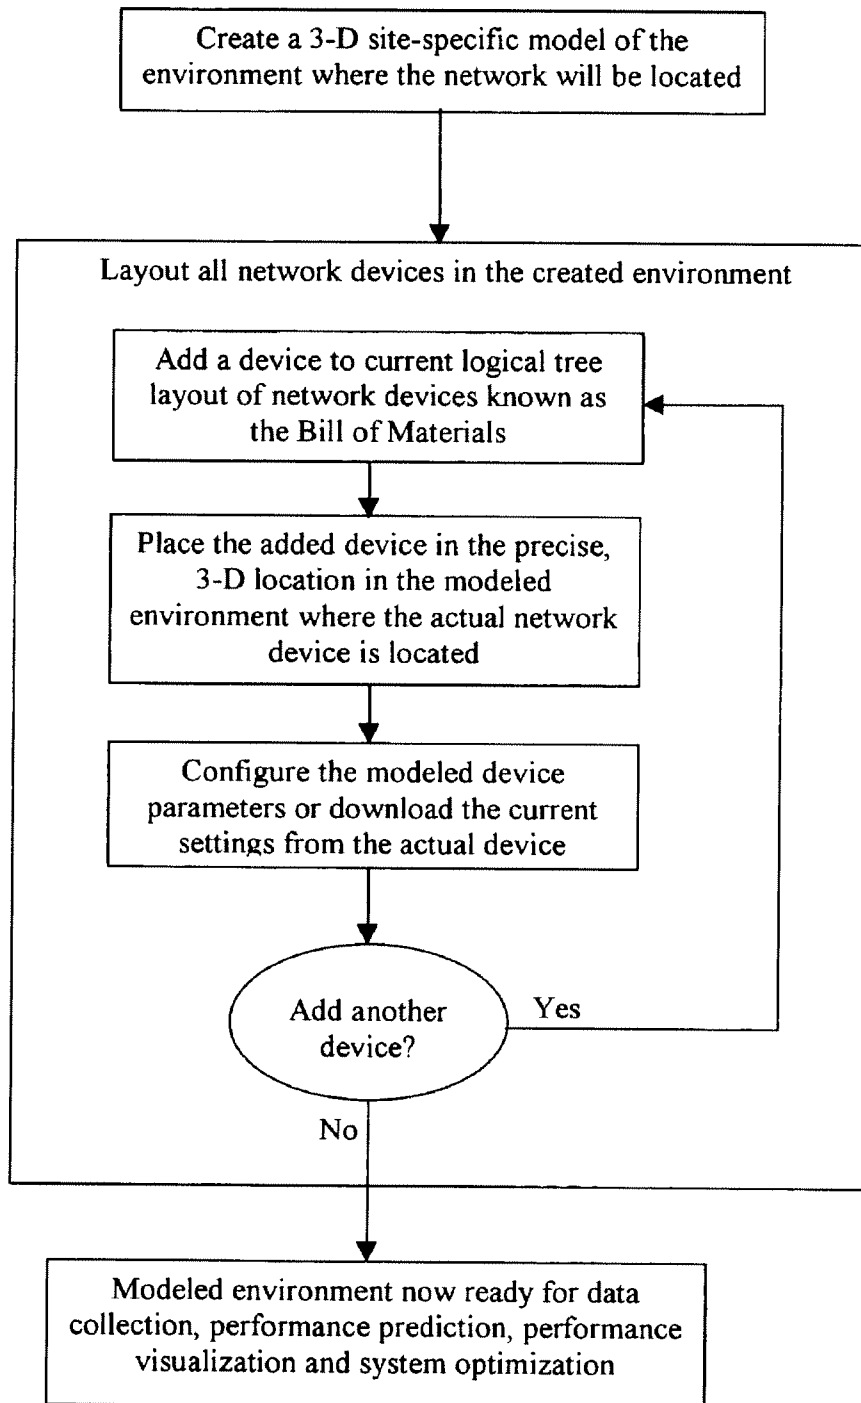
U.S. Patent

Dec. 6, 2005

Sheet 3 of 6

US 6,973,622 B1

Figure 5: Method for creating a 3-D site specific model of the environment



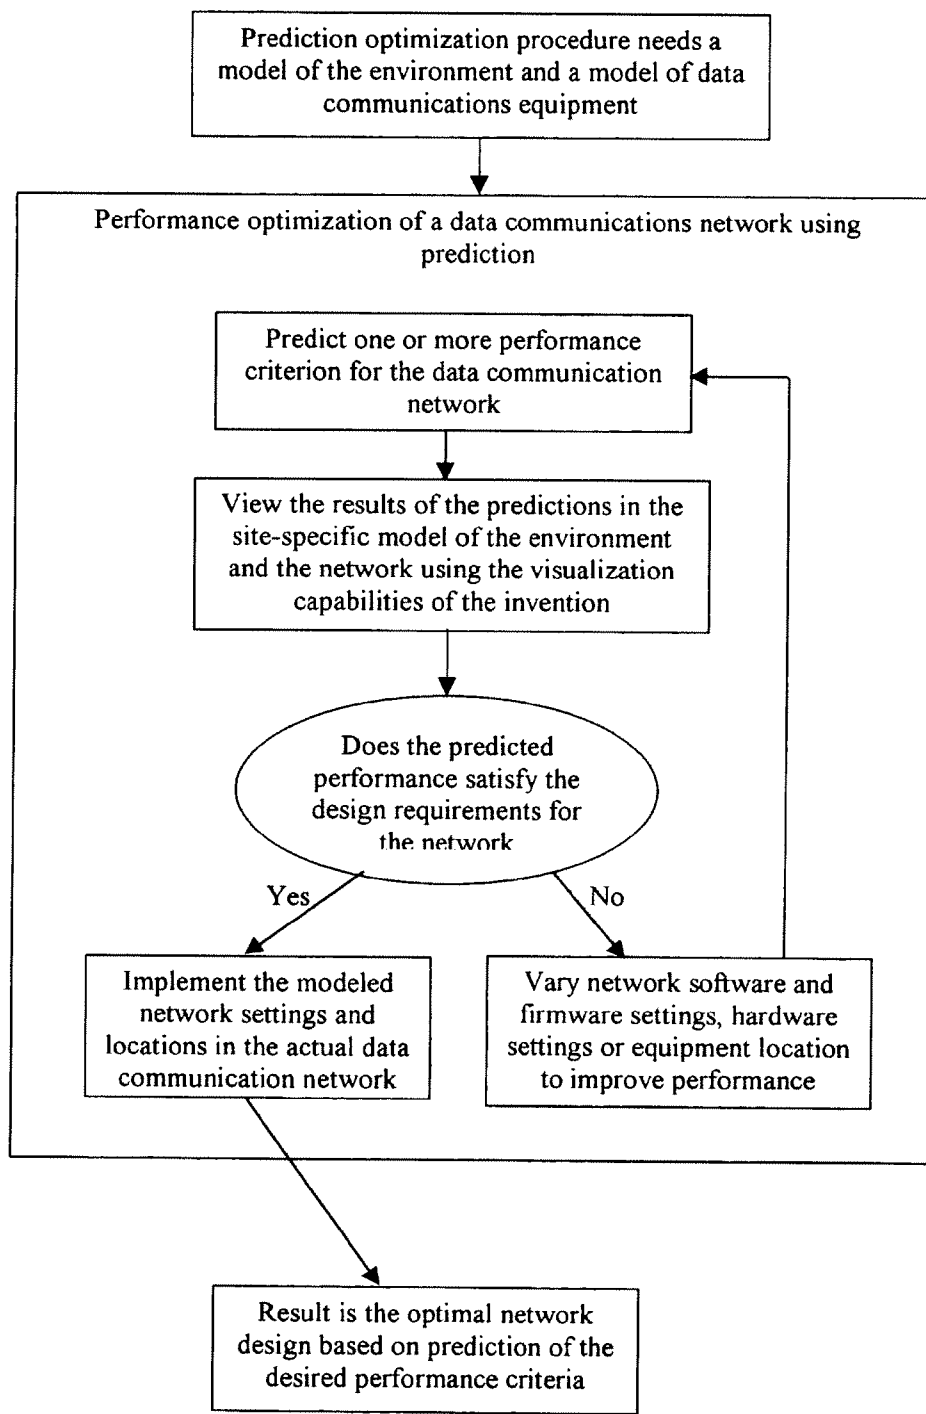
U.S. Patent

Dec. 6, 2005

Sheet 4 of 6

US 6,973,622 B1

Figure 6: Method for optimizing a data communications network using predictions



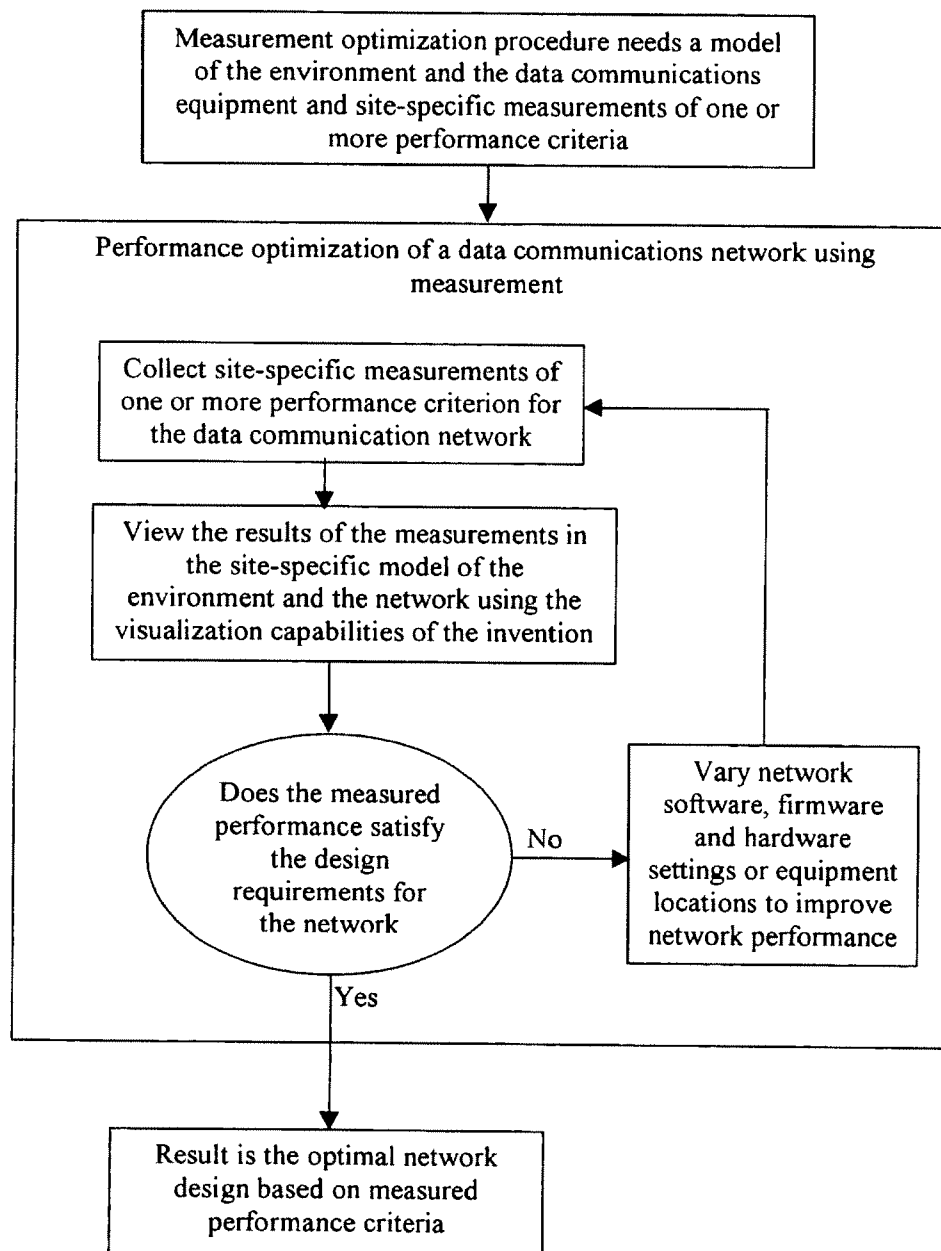
U.S. Patent

Dec. 6, 2005

Sheet 5 of 6

US 6,973,622 B1

Figure 7: Method for optimizing a data communications network using measurements



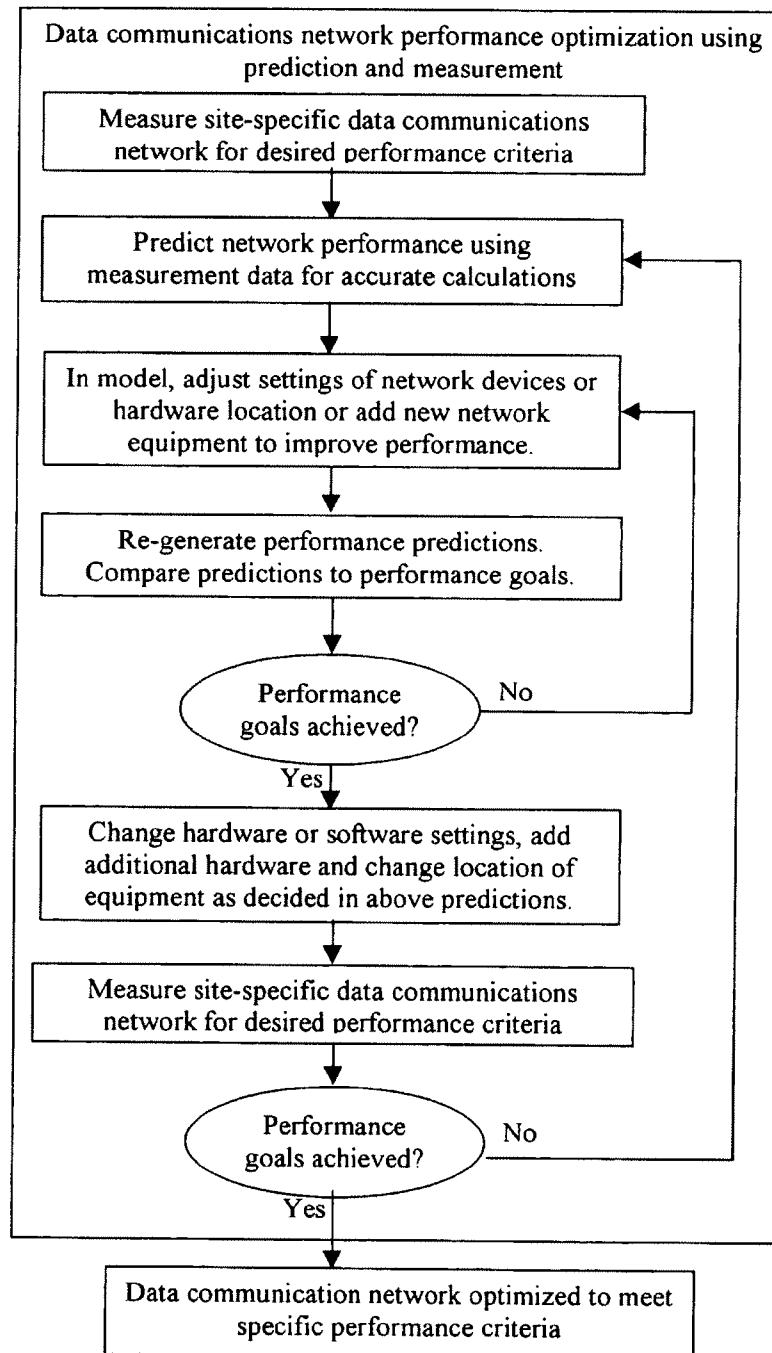
U.S. Patent

Dec. 6, 2005

Sheet 6 of 6

US 6,973,622 B1

Figure8: Method for optimizing a data communications network using predictions and measurements.



US 6,973,622 B1

1

SYSTEM AND METHOD FOR DESIGN, TRACKING, MEASUREMENT, PREDICTION AND OPTIMIZATION OF DATA COMMUNICATION NETWORKS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to application Ser. No. 09/318,842, entitled "Method and System for Managing a Real Time Bill of Materials," filed by T. S. Rappaport and R. R. Skidmore, now U.S. Pat. No. 6,493,679, Ser. No. 09/318,841, entitled "Method And System for a Building Database Manipulator," filed by T. S. Rappaport and R. R. Skidmore now U.S. Pat. No. 6,850,946, Ser. No. 09/318,840, entitled "Method and System For Automated Optimization of Communication component Position in 3D" filed by T. S. Rappaport and R. R. Skidmore, now U.S. Pat. No. 6,317,599. Pending application entitled "Method and System for Designing or Deploying a Communications Network which Allows Simultaneous Selection of Multiple Components" filed by T. S. Rappaport and R. R. Skidmore, Ser. No. 09/633,122, filed on Aug. 4, 2000, as well as applications entitled "Method and System for Designing or Deploying a Communications Network which Considers Frequency Dependent Effects", Ser. No. 09/632,121, filed by T. S. Rappaport and R. R. Skidmore on Aug. 4, 2000 now U.S. Pat. No. 6,625,454, as pending application entitled "Method and System for Designing or Deploying a Communications Network which Considers Component Attributes", Ser. No. 09/632,853, filed by T. S. Rappaport, R. R. Skidmore, and Eric Reifsnider on Aug. 4, 2000, as well as application entitled "Improved Method and System for a Building Database Manipulator", Ser. No. 09/633,120, filed by T. S. Rappaport and R. R. Skidmore, now U.S. Pat. No. 6,721,769 and pending application entitled "System and Method for Efficiently Visualizing and Comparing Communication Network System Performance", Ser. No. 09/632,803 filed by T. S. Rappaport, R. R. Skidmore, and Brian Gold on Aug. 4, 2000, and co-pending application "Method and System for Automated Selection of Optimal Communication Network Equipment Model, Position and Configuration in 3-D", Ser. No. 09/667,689, filed by T. S. Rappaport, R. R. Skidmore, and P. SheethaNath filed concurrently, the subject matter of which is incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The invention relates to the field of communications networks, and more specifically to the design thereof, and the measurement, visualization, prediction and optimization of the performance of data communication networks. A method and system to predict, visualize and optimize the performance of data communication networks is used to design, measure, monitor, troubleshoot and improve these data networks using an accurate site-specific model of the physical environment and the components comprising the data network.

2. Description of the Related Art

Communications networks are used to send information from one place to another. This information often takes the form of voice, video or data. To transmit information a communications network breaks down a message into a series of numbers. These numbers describe how to construct the information using some predetermined method. For example, the numbers could represent digital samples of the

2

signal voltage that should be applied to a speaker so that the speaker reproduces the sound of the voice, as shown in FIG. 1. The information is in this case the voice message, which was transmitted over the communications network.

The process of representing information can be analog or digital. In an analog communications network the message that is transmitted is a continuously changing number. In a digital network, numbers that change at discrete, regular intervals, instead of continuously represents the message. The signal is represented by a single number each interval. This number may be converted to a binary form so that the entire message can be represented as a finite number of ones and zeros. Each binary digit in the message is called a bit. These bits are transmitted and interpreted by the receiver as the message. Binary and digital versions of a signal are shown in FIG. 2.

Data communication networks are a specific type of communication network that transmit digital information, represented as bits or bytes (a group of 8 bits), in an indoor or outdoor, wired or wireless network from a transmitter to a receiver. While conceptually simple, the means of transmitting the data from some point A to some point B are complicated and varied in implementation. Hundreds of protocols, hardware devices, software techniques and programs exist to handle how data is sent correctly and efficiently. The exact performance of a given data communication network is extremely difficult to predict or even measure because of this complexity and additionally because of the performance effects of the time varying nature of data communications networks and the channels they operate in.

Data communication network can be classified as either a circuit switched or a packet switched network. Both network types use channels to transmit information. A channel is a named communications path between users of a communications network. A channel may consist of many different individual hardware devices and is a specific route between a transmitter and a receiver. In a circuit switched network, information is transmitted by way of an exclusively reserved channel. A network channel is reserved for the sole use of a single transmission and bits are sent all at once. An example of this is the transmission of a document using a fax machine. In this case the fax machine converts the image of the document into pixels. Each pixel is a small, dot-sized, rectangular piece of the paper. Each pixel is considered to be either black or white. The data that will be transmitted is a series of bits that represent whether each dot is black or white. When the message (in this case an image of a document) is ready to be sent from one fax machine to another, a telephone circuit is dedicated to the data transfer by placing a telephone call on the plain old telephone system (POTS) communications network. The telephone line is used exclusively by the fax transmission, making it a circuit switched transmission. After establishing a connection, all data is sent from the first fax machine to the second in a single, long stream of bits. The bits in this case are transmitted as different frequency tones on the telephone line. A high pitched tone may represent a "1" while a low pitched tone may represent a "0." The receiving fax receives the bits of the message by translating the series of high and low pitch tones into data bits. The receiving fax machine will then be able to reconstruct a copy of the original document by drawing a black dot at the locations indicated by the data bits.

Packet switched networks are another type of data communication networks in which all data bits are transmitted as many, small chunks of data bits called packets and sent individually from one location to another. A packet is a

US 6,973,622 B1

3

self-contained portion of a full message that is made up of a header, data bits, and sometimes footer. The packet contains information in the header and footer that allows the data communications network to properly transmit the packet and to know of which message the data in the packet is a part. The header generally is labeled with an identifier that the network uses to forward the packet to the correct receiver. The header and footer information are often used to reassemble the packet with other packets to reform the original message and to check if errors were made in the transmission of the packet. The receiver can assemble all received packets into the original message by throwing away the header and footer headings and reassembling the data bits from all packets into the original message.

Packet switched networks are classified as connection oriented or connectionless depending on how the packets are transferred. In connection-oriented networks, a network channel is used predefined for each transmission. While this transmission can consist of multiple packets, the route from transmitter to receiver is already established, so that all packets sent on this channel can immediately be sent directly to the receiver. Whereas, in connectionless networks, packets are sent simultaneously on a shared channel in multiple transmissions. In this case, packets require an identifier that gives the address of the receiver. This address is understood by the communications network to allow the packet to be properly sent to the correct receiver. Since each packet can be transmitted separately and thus interleaved in time with packets from other transmissions, it is generally more efficient to use a connectionless transmission method when using shared network resources.

An example of a connectionless, packet-based transmission is a file transfer between two computers on an internet protocol (IP) based, Ethernet network that both computers are attached to. In this case, the file that is to be transmitted is fragmented at the transmitter into appropriate packets and labeled with the IP address, which is the identifier used by the network to forward the packet to the correct receiver. The packets are then sent from the transmitting computer to the receiving computer. The Ethernet network is capable of supporting multiple file transfers from many different computers all using the same network by controlling the flow of packets from each destination in a shared fashion. The receiver then assembles the packets into an exact copy of the original file, completing the transmission.

All data networks utilize some form of communication protocol to regulate the transmission and reception of information. A protocol is the set of rules that all hardware and software on a communication network must follow to allow proper communication of data to take place. Many hundreds of protocols are in active use today in the worldwide exchange of information. Some of these protocols, such as the Transport Control Protocol (TCP) or the User Datagram Protocol (UDP), define the way in which the network is accessed. Other protocols, such as the Internet Protocol (IP) or the File Transfer Protocol (FTP), define how messages and packets are formatted, transmitted, and received.

All data communication networks may be analyzed in some fashion to evaluate the efficiency and performance of the network as well as to confirm the network is functioning properly. In order to evaluate the functionality of these data networks, certain performance criterion is used. These performance criteria include, but are not limited to: throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter,

4

bandwidth delay product and handoff delay time. Each performance criterion specifies a different performance parameter of a data communications network. These criteria are further described below.

A link is a portion of a path followed by a message between a transmitter and a receiver in a data communications network. Network connection often consists of individual devices relaying network packets from the transmitter to the receiver. This means a network connection can consist of several actual transmissions between the original transmitter and the intended receiver. Each individual relay is called a link. Typically a full network connection consists of several links. Performance criteria can be measured for each individual link.

Throughput is a measurement of the amount of data, which can be transmitted between two locations in a data network, not including header, footer or routing information bits. It is generally measured in bits per second (bps) and can be specified for hardware, software, firmware or any combination thereof that make up a connection between transmitter and receiver in a data communication network. Bandwidth is similar to throughput as it is defined for data communication networks. Bandwidth is the raw data rate that may be sustained by a given communications network and is generally slightly higher than throughput. For instance, an Ethernet link may be rated for a 10 Mbps bandwidth but a measurement of an actual file transfer may show that the rate at which data can actually be transferred between two computers using that same link is only a throughput of 6.8 Mbps as is taught in Peterson, L. L. and Davie, B. S., *Computer Networks: A Systems Approach*. San Francisco: Morgan Kaufmann Publishers, 2000.

Quality of service (QoS) is a term that is used to describe networks that allocate a certain amount of bandwidth to a particular network transmitter. Such a network will allow a transmission to request a certain bandwidth. The network will then decide if it can guarantee that bandwidth or not. The result is that network programs have a reliable bandwidth that can more easily be adapted to. When the quality of service of a connection is measured, the bandwidth that the network claims to offer should be compared to the actual bandwidth for different requested bandwidths.

FIG. 3 illustrates the difference between bits, packets, and frames. Various error rates are defined for data communication networks for bits, packets and frames. Bits are the core of packets and frames. The bits are the actual message data that is sent on the communications network. Packets include the data bits and the packet header and packet footer. The packet header and packet footer are added by communications network protocols and are used to ensure the data bits are sent to the right location in the communications network and interpreted correctly by the receiver. The packet header and packet footer are also used to ensure that packets are sent correctly and that errors are detected should they occur. Frames are simply series of bits with a certain pattern or format that allows a receiver to know when one frame begins or ends. A bit error rate is the percentage of bits that reach the receiver incorrectly or do not reach the receiver as compared to the number of bits sent. Packet error rate or dropped packet rate is the percentage of packets that reach the receiver incorrectly or do not reach the receiver as compared to the number of packets sent. A frame error rate is the percentage of frames that reach the receiver incorrectly or do not reach the receiver as compared to the number of packets sent.

Several terms are used to quantify the delay times of certain network events and may be expressed in time units

US 6,973,622 B1

5

of seconds. Packet latency is the time required to send a packet from transmitter to receiver, while Round Trip Time (RTT) is the time required for a packet to be sent from transmitter to receiver and for some sort of acknowledgment to be returned from the receiver to the original transmitter. Propagation delay, transmission delay, processing delay, and queuing delay describe the time required for different portions of a packet transmission to occur. The packet latency and round trip time of a network connection is found by summing the propagation delay, transmission delay, processing delay and queuing delay of either a one way or round trip network connection. Propagation delay is the time required for a packet to traverse a physical distance from the transmitter to the receiver. Transmission delay is the time required from when the first bit of a packet arrives for the last bit of the same packet to arrive. Processing delay refers to the time required to subdivide a data message into the individual packets at the transmitter, and to the time required to recreate the full data message from the data packets at the receiver. Queuing delay refers to the time spent waiting for shared resources to be freed from use by other transmissions. These delay times are all useful for evaluating different aspects of a data communications network performance.

Two other network performance criteria are packet jitter and bandwidth delay product. Packet jitter is the variation in the arrival time of packets that are expected to arrive at a regular rate and is typically measured in time units of seconds. A bandwidth delay product is the number of bits that can be sent from a transmitter before the first bit sent actually reached the receiver. The bandwidth delay product is found by multiplying the packet latency of a certain link by the bandwidth of the same link.

Handoffs occur in wireless data networks when a user moves out of range of one access point and into range of another access point. In this situation, the first access point must pass the responsibility of delivering data to the wireless user to the second access point. The handoff time is the amount of time required by an access point to coordinate with another access point to allow a wireless user to connect from one access point to another access point.

Software utilities and hardware devices have been developed to measure the performance statistics of data communication networks throughout the lifetime of data communication networks. Some of the more common and relevant tools are briefly described here.

A large number of command line tools are available to quickly allow a computer user to measure the approximate network performance a connection. Many command line programs are widely used on Windows, UNIX, and Macintosh operating systems and are somewhat useful for diagnostic and troubleshooting work on data networks. Examples of these command line programs include ping and traceroute. Using the ping command line program, it is possible to measure approximate data latency between different data network devices and confirm that a network connection is available between the two devices. Network connections often consist of individual devices relaying network packets from the transmitter to the receiver. This means a network connection can consist of several actual transmissions between the original transmitter and the intended receiver. Each individual relay is called a link. Typically a full network connection consists of several links. Thus, using traceroute, a probable path from relaying device to relaying device between the transmitter and the receiver can be determined so that the exact links used by the network transmissions are known. Additionally, using trac-

6

eroute, the time required to traverse each individual link can be measured, and individual links that may not be functioning properly can be identified.

Various command line tools that are not included with operating systems have also been developed for somewhat more accurate, though still approximate, network measurement tasks. Some examples of these tools include ttcp, and tcpdump. ttcp stands for Test TCP <http://www.pcausa.com/Utilities/pcattcp.htm> and is a free utility originally written for the BSD Linux operating system, but is now available for other UNIX operating systems as well as Microsoft Windows. ttcp is a basic point-to-point throughput measurement program that allows the user to control buffer sizes, various low level TCP or UDP options and control the exact data that is sent.

tcpdump is a simple utility from the class of tools called packet sniffers. Packet sniffers allow a network administrator to view the content, including header and footer information, of actual packets on a network. tcpdump allows a user to view (or "sniff") packets that are received by a host (though not necessarily intended for that host) and display all headers that match a certain user configurable pattern. tcpdump is a useful tool for troubleshooting network connections because it allows the user a direct view of the exact network traffic.

Pathchar is a UNIX command line utility which is capable of measuring the throughput between each network relay device (e.g. a router, hub or switch) in a data communications network by varying the size of the test packets that it transmits and measuring the latency of that packet transmission to various network points. The tool functions very similarly to traceroute but adds the ability to measure throughput (albeit indirectly), not just latency. Pathchar is only limited by the network hardware in the links it measures. The program needs a hub, switch or computer to transmit an acknowledgement to the test packets. This means that hidden links that do not transmit acknowledgements such as Ethernet bridges can not be measured individually by pathchar.

Several companies produce network measurement, monitoring, tracking and forecasting utilities. Some of the commonly used utilities are discussed below. The tools selected are illustrative of the state of the art of network performance measurement and asset tracking.

netViz, made by netViz Corporation, is a visual database program that allows a network administrator to track network equipment in terms of its physical location and in terms of its logical layout. This program allows the user to input the settings, locations, and configurations of the network and track the assets in your network. The tool is capable of storing this data in a two dimensional geographic map or floor plan of a building, but can not track devices in a three dimensional manner. The tool, also, does not provide network testing, measurement or monitoring features, nor does it support communication prediction or performance visualization capabilities for data communication networks. It is simply a database for accurate and useful tracking of assets.

NetIQ Corporation (was Ganymede Software, Inc.) makes a network monitoring and forecasting tool called Chariot. Chariot is able to measure throughput and many other network statistics for all popular network types, operating systems and protocols available today. The program uses a server and several small agent programs to collect data. The server checks each agent, installed on user's computers throughout the network, at regular intervals and uses them to measure network characteristics while storing

US 6,973,622 B1

7

the results on the server. These agents can measure the network connection to the server or to one another and are capable of simulating the traffic patterns of any network program and any desired usage pattern of one or more hypothetical users. The program is also capable of using the measured data to forecast expected network traffic and conditions.

Visonael Corporation (was NetSuite Development Corporation) makes several network tracking and measurement products, including NetSuite Audit, Design and Advisor. These software products are capable of automatically detecting the network equipment in use. This information as well as manually entered information can then be placed in a physical or logical diagram of the network. Visonael also offers a product to verify that networks have been configured properly and can make recommendations for configuration changes and upgrades to your network. The software products are unable to predict or measure the performance in a site-specific manner and are not capable of predicting the performance of wireless based data communication networks.

SAFCO Technologies, Inc. (now a part of Agilent Technologies) has recently created several wireless data measurement and prediction products. SAFCO makes a product called DataPrint, which is used to measure various data performance parameters of mobile telephone data networks. Their WIZARDS product also supports analysis of the effects of wireless data transmission on the overall capacity and Quality of Service for a wireless telephone network.

Wireless Valley Communications, Inc. has created a new concept called SitePlanner, which is capable of measuring and tracking the site-specific network performance of a data communications network in a physically accurate three-dimensional model of an environment. SitePlanner uses a software module called LANFielder to measure throughput, packet latency and packet error rates for any wired or wireless network connection in any Internet Protocol (IP) data communications network. Additionally, SitePlanner allows a full network to be modeled in a physically accurate manner so that precise measurements and performance predictions can be made in a site specific way. SitePlanner also allows a logical layout of a network to be stored simultaneously with a physical layout. The tool also stores both a logical interconnection and a site-specific model of any communications network using a Bill of Materials format.

In addition to network measurement and asset management tools, a good deal of research has taken place in the field of wireless data communication network performance. The research described below represent the work, which pertains to the field of this invention.

Xylomenos and Polyzos have explored the performance of UDP and TCP packets sent over several fixed, IEEE 802.11 wireless LAN network connections in Xylomenos, G., Polyzos, G. C. "TCP and UDP Performance over a Wireless LAN" *Proceedings of IEEE INFOCOM*, 1999. The research has focused on throughput limitations caused by software implementation issues and operating system shortcomings. The researchers used their own modified version of the command line utilities *tcpdump* and *netstat* under Linux to perform UDP and TCP throughput tests. All measurements were taken between three fixed locations and focused on varying the wireless LAN card types (PCMCIA or ISA) and the end-user computer hardware (i.e. Pentium 150 with 48 MB of RAM vs a Pentium 200 MMX with 64 MB of RAM). The conclusions the researchers make are recommendations for changes in the implementation of

8

network protocols and linux operating system enhancements. The measurements did not consider the effects of different physical locations or the effect of variations in the wireless communications channel on the network throughput.

Maeda, Takaya and Kuwabara have published a measurement of wireless LAN performance and the validity of a Ray tracing technique to predict the performance of a wireless LAN network (Maeda, Y., Takaya, K., and Kuwabara, N., "Experimental Investigation of Propagation Characteristics of 2.4 GHz ISM-Band Wireless LAN in Various Indoor Environments," *IEICE Transactions in Communications*, Vol. E82-B, No. 10 Oct. 1999). The measurements were tracked in a small, highly radio frequency (RF) controlled environment and indicated that the wireless LAN throughput and BER were correlated to the delay spread of the wireless channel. The researchers have not however presented any way to actually predict a bit error rate or throughput from the predicted delay spread profile output by a ray tracing technique.

Duchamp and Reynolds have presented IEEE 802.11 wireless LAN, packet throughput measurement results for varying distances in Duchamp, D., and Reynolds, N. F., "Measured Performance of a Wireless LAN," *Local Computer Networks*, 1992. *Proceedings, 17th Conference on*, 1992. These measurements were performed in a single hallway. Thus, these measurements, too, suffer from failing to measure a representative environment. The researchers did not present a model to predict their results nor did they attempt to validate any sort of computer prediction technique.

Bing has also presented measured results of the performance of IEEE 802.11 Wireless LAN in "Measured Performance of the IEEE 802.11 Wireless LAN," *Local Computer Networks*, 1999. *LCN '99. Conference on*, 1999. Bing presents delay and throughput measurements as well as theoretically based throughput and delay time tabulations for various wireless LAN configurations. The results are given as optimal results, however. All measurements were performed in such a way that the wireless channel had the least possible effect on the overall throughput and delay times. Therefore, the results presented are an upper bound on best possible results and do not extend into a site-specific wireless LAN performance prediction technique.

Hope and Linge have used measurements to calculate the needed parameters for predicting the coverage area of a Wireless LAN network in an outdoor environment by using the Okumura model. The researchers have made outdoor measurements with standard IEEE 802.11 wireless LAN modems to calculate the needed parameters of the Okumura model and have presented these results in Hope, M. and Linge, N., "Determining the Propagation Range of IEEE 802.11 Radio LAN's for Outdoor Applications," *Local Computer Networks*, 1999. *LCN '99. Conference on*, 1999. Using these results, The coverage area outdoors could be calculated. However, the results do not allow the user to predict the performance in terms of throughput or latency of a wireless LAN.

Several patents related to, and which allow, the present invention are listed below:

U.S. Pat. No. 5,491,644 entitled "Cell Engineering Tool and Methods" filed by L. W. Pickering et al;

U.S. Pat. No. 5,561,841 entitled "Method and Apparatus for Planning a Cellular Radio Network by Creating a Model on a Digital Map Adding Properties and Optimizing Parameters, Based on Statistical Simulation Results" filed by O. Markus;

US 6,973,622 B1

9

U.S. Pat. No. 5,794,128 entitled "Apparatus and Processes for Realistic Simulation of Wireless Information Transport Systems" filed by K. H. Brockel et al;
 U.S. Pat. No. 5,949,988 entitled "Prediction System for RF Power Distribution" filed by F. Feisullin et al;
 U.S. Pat. No. 5,987,328 entitled "Method and Device for Placement of Transmitters in Wireless Networks" filed by A. Ephremides and D. Stamatelos;
 U.S. Pat. No. 5,598,532 entitled "Method and Apparatus for Optimizing Computer Networks" filed by M. Liron et al.
 U.S. Pat. No. 5,953,669 entitled "Method and Apparatus for Predicting Signal Characteristics in a Wireless Communication System" filed by G. Stratis et al.
 U.S. Pat. No. 6,061,722 entitled "Assessing Network Performance without Interference with Normal Network Operations" filed by W. J. Lipa et al.
 U.S. Pat. No. 5,831,610 entitled "Designing Networks" filed by D. L. Tonelli et al.
 U.S. Pat. No. 5,821,937 entitled "Computer Method for Updating a Network Design" filed by Tonelli et al.
 U.S. Pat. No. 5,878,328 entitled "Method and Apparatus for Wireless Communication System Organization" filed by K. K. Chawla et al.

An existing product, SitePlanner, described in patent application Ser. Nos. 09/352,678, 09/221,985, 09/318,842, 09/318,841, 09/318,840, and other inventions cited previously, are useful for designing, measuring and optimizing communication networks because the products can predict radio frequency effects directly relevant to any communication network for any physical location. That is, using information about the physical layout of any communications network and the configuration of its hardware, prior art can provide a visual display of the expected received signal strength intensity (RSSI), signal to noise ratio (SNR), relative received power intensity, best server, and equal power location, as well as other useful parameters for voice and data networks, for any modeled physical location. These statistics can be predicted for the forward link (from a transmitter to a receiver), or for the reverse link (replies from the original receiver to an original transmitter) directions for wireless networks. The site-specific nature of these predictions translates directly into quick and useful visualizations of the quality of a communication network. However, the prior art does not consider methods for properly modeling (e.g. predicting) the complexities that go into determining the values for actual network operating performance parameters that are simultaneously affected by multipath propagation, multiple interfering data transmissions from multiple sources, signaling protocols, equalization methods, and the like. Predicting bit error rates, data throughput, delay, and quality of service metrics in a 3-D physical model of an actual site-specific environment is a very difficult task, and one which has not been solved heretofore, since different modem vendors have different and often-times proprietary methods for mitigating or dealing with multipath, multiple access interference, protocol type, packet size, and noise. That is, the state of the art shows how to measure and display and make predictions for basic communication metrics but does not provide specific prediction algorithms for a wide range of important data network performance parameters in a reliable, site-specific manner. Simply put, a wireless network performance prediction engine, which is able to consider an accurately modeled 3-D physical environment, and which exploits knowledge of specific component layouts, is not found in the prior art and is not obvious due to the complex nature of having to account for all possible physical, electrical, and logical factors for all components in a

10

network, as well as the factors within the channel of a wired or wireless network, that lead to actual network performance.

Prior published papers in the area of communications networks do not demonstrate the ability of any invention to accurately predict three dimensional, site-specific network performance criteria. The paper mentioned earlier by Maeda, Y., Takaya, K., and Kuwabara, N., "Experimental Investigation of Propagation Characteristics of 2.4 GHz ISM-Band Wireless LAN in Various Indoor Environments," *IEICE Transactions in Communications*, Vol. E82-B, No. 10 Oct. 1999 has demonstrated the ability to predict the delay spread of a wireless channel and that the prediction correlates well with throughput, but the described method is not actually able to predict throughput or any other network performance criteria. While some prior art has demonstrated the ability to track network assets in a two dimensional manner with some physical accuracy, these products have not contemplated the ability to predict future network performance for similar or different physical environments (e.g. installations). Many products allow the measurement of network performance criteria, but no prior art has contemplated a 3-D representation of the physical environment with the physical installed base of components, for the purpose of predicting network performance parameters. Furthermore, no tool or invention exists that can directly measure, track the assets of, predict the network performance criteria of, and visualize the network performance criteria of a data communications network in a three-dimensional site-specific manner.

Furthermore, none of the prior art has considered an invention that can perform precise, site-specific, three dimensional performance prediction of complicated network parameters using a priori measurements from an existing network, or by using the site-specific layout details of particular components within a data communications network. Furthermore, none of the prior art has autonomously measured site-specific network performance parameters from an actual network system or subsystem using a system of agents, and then applying the specific 3-D locations and measured results of those measurement agents to create a 3-D prediction model for future network performance in the same, similar, or different physical environments. Furthermore, none of the prior art has developed a hierarchical system of measurement and prediction engines, that have the ability to measure network performance parameters in the field and have the ability to produce a predictive engine for network performance parameters that can be shared with remote prediction engines, for the purpose of measuring and predicting network performance in a 3-D site-specific manner.

The present invention extends the prior art in a non-obvious way to provide wireless and wired network performance prediction, visualization and measurement for important data communications-specific performance criteria, also called performance parameters, such as throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time in a site-specific, three dimensionally accurate manner. The invention contemplated here allows novel distributed measurement techniques for the above performance parameters. Furthermore, prediction methods for the above performance parameters are created, which use network measurements or applied values derived from other means, and which also use the

US 6,973,622 B1

11

radio frequency environment, the 3-D physical network layout, the channel propagation characteristics of a site-specific environment, and the specific physical layout of components, for the computation of predicted performance parameter values.

SUMMARY OF THE INVENTION

The present invention is capable of predicting, measuring, and optimizing the performance of a data communications network. The invention is capable of representing a detailed layout of a fully deployed or contemplated communications network within a physically accurate computer representation or model of a three dimensional environment. This allows the invention to store measurements and determine performance predictions within a site-specific representation of the physical environment, while using specific information about the network entities, components, subsystems, and systems used to create the actual or contemplated network. Measurement agents, with known or assigned 3-D position locations, are used to measure in-situ performance parameters that are transmitted to a server processor. The server processor has an accurate 3-D model of the environment, and is able to process the measured data, and is also able to provide predictive models using site-specific information that may be independent of or may make use of measured data. The server process is able to communicate with other server processors in a hierarchical manner, such that data fusion from many remote or collocated networks may be assembled and used for display and cataloging of measurements that may or may not be used for creation of predictive performance models. Alternatively, each server processor is able to compute predictive performance models without the use of measured data, by simply considering the site-specific layout of physical components, as well as the specific delay times, transit times, propagation effects, and multipath and noise factors within the physical network.

The invention can predict throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time in a site-specific, three dimensional model of any environment. The invention can measure and predict all of the above performance criteria and store the results in the physically accurate three-dimensional model of a data communications network and the environment in which it is installed. Further, the invention can display the measured and predicted performance criteria for any data communications network in the three dimensions, site-specific model of the environment. These capabilities provide a powerful design environment for wired and wireless networks, which allows one skilled in the art to quickly and easily design, measure, predict, optimize and visualize data network communication performance criteria in a three dimensional, site-specific manner using methods never before contemplated.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1: Example transmission of data over a communications network

FIG. 2: Creation of a digital signal from an analog signal

FIG. 3: Illustration of the difference between bits, packets and frames.

FIG. 4: Illustration of the data displayed in each node of the Tree View of a data communications network.

12

FIG. 5: Method for creating a 3-D site-specific model of the environment

FIG. 6: Method for optimizing a data communications network using predictions

FIG. 7: Method for optimizing a data communications network using measurements

FIG. 8: Method for optimizing a data communications network using predictions and measurements.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS OF THE INVENTION

The present invention contemplates the abilities to design, measure, predict and optimize the performance of a data communication networks. The invention uses an accurate computer generated three-dimensional model of a communications network stored in a computer database environment. The invention allows the user to place the network cables, hubs, routers, switches, bridges, wireless access points, amplifiers, splitters, antennas (point, omnidirectional, directional, leaky feeder, distributed, array, etc.) transceivers, terminators and other communications and computer networking equipment in their actual modeled physical locations. The present invention uses this highly accurate model of the physical layout of infrastructure to allow a user to visualize, predict and optimize the performance of any communication network in any 3-D site specifically modeled physical location.

The present embodiment of the invention is capable of modeling the site-specific communications network hardware from both a logical connection and a physical location perspective. The invention uses well-known hierarchical, logical connection concepts (sometimes called topological layout) suited for data communications networks in combination with a physically accurate, site-specific model of the data communications network. Previous inventions focus on only the topological, or relational, layout of network components with one another. This invention uses specific 3-D modeling and, therefore, allows highly accurate asset management and facilities tracking of actual installed equipment while simultaneously providing for network performance prediction, measurement, and design capabilities that exploit the exact physical dimensioning of the network. In addition, the invention simultaneously stores an inventory of important network-specific and equipment-specific characterizations of all objects used in the network, such as vendor, model number, network hardware type, operating system version, firmware and software type and version. The hierarchical, tree based model of the network is termed the Layout View. The physically accurate, site-specific model of the network is termed the Site View, whereby the attributes of each device can be displayed, stored or printed by selecting a particular item or node within the 3-D environmental model. Further, network hardware and software components can be interactively replaced, removed, reconfigured or moved to a new location in real-time using either the Layout View or the Site View. Each of these ways of tracking and designing a network in a 3-D site specific model of the environment with accurate dimensioning of true spatial position are further described below and are used to create a Bill of Materials for the modeled data communications network, whereby a preferred embodiment is described in co-pending patent application "Method and system for designing or deploying a communications network which considers component attributes," filed on Aug. 4, 2000.

US 6,973,622 B1

13

An example of some of the information contained in the Layout View, hierarchical layout of a data communications network is shown in FIG. 4. In the figure, a tree structure is used to display all hardware in the network. Each node in the tree contains information which is used to track the true physical location, logical layout and electrical, optical and electromagnetic connections for the data communications network hardware as well as any version numbers and settings of software or firmware running on that network equipment and the known performance parameters of that equipment, including the device throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time.

The Site View of the invention has a physically accurate, three-dimensional modeling capability to display all network devices in a site-specific model of the environment that the network is located in. That is, the preferred embodiment of the invention allows each modeled hardware and software device to be placed in a three-dimensionally accurate manner and to track attributes of that device relevant to data communications networks. These key attributes include such items as the hardware type, hardware configuration, software type, software configuration, operating system version, as well as upper, lower and "typical" specifications for each component. These specifications may include important device or network subsystem operating parameters, such as throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time. As described below, the Site View supercedes prior art described in previous co-pending patent applications by Wireless Valley Communications, Inc by hereby considering the difficulties and solving data network prediction, design and optimization problems for more complicated data communication networks. Specifically, this new invention considers physical, site-specific modeling techniques and performance prediction methods and design methods for data network systems, both wired and wireless, which have performance characteristics that are based on much more complicated physical factors than just radio signal strength, interference, or multipath alone. In particular, for data communication networks, many additional factors, which relate to particular network equipment or modem designs, such as packet size, equalizer deployment, modulation methodology, source and error coding methods, packet protocols, as well as the number of co-channel network users, the type of persistency used for packet retransmission, or the multipath propagation effects in a wireless system, provide additional factors that must be considered in the design of a communication network that is designed for data traffic as opposed to simply voice traffic.

One difficulty that today's network designer or network system administrator faces is that most networking equipment uses proprietary, non-public methods for implementing various network devices, and these methods vary by specific vendor. Thus, it is difficult to form reliable prediction models by just using basic physical propagation models in a wireless network, for example. As data transmission technologies such as Bluetooth, DSL, Voice over IP, and future packet-based cellular radio network architectures proliferate, the ability to predict and measure specific network performance parameters will become increasingly important, and the

14

ability to properly incorporate measurements into 3-D prediction models for performance parameters will be important for proper network deployment.

This invention considers attributes relevant to packet-switched data communication networks, which require more extensive and non-obvious modeling when compared to traditional cell phone or telephone voice communication systems that are circuit switched and use a dedicated single user (or bounded number of users) per assigned operating channel. Data communication networks have performance criteria that are specific to packet-based systems and that are not useful to all types of communication networks contemplated previously. For this reason, the preferred embodiment of the invention can additionally predict the throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time, based on the specific physical and spatial location of each network component, as well as the physical, electrical, and logical attributes of the specific components. The performance prediction methods take into account all devices and network equipment, including the physical locations within the 3-D modeled environment, using the constructed Bill of Materials of the network within the 3-D modeled environment, and is capable of performance predictions for any desired location in the modeled network and environment, where a location may be within a room, at a particular location in a room, within a building, or in an outdoor region of varying granularity, depending on the requirements of the user.

Prediction of throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time and other performance parameters may be carried out by predicting the performance for all wired network components separately from the performance of wireless components, and then combining the results to get the net network performance. To predict the performance of a wired communication link, it is important to combine the known effects of each piece of wired equipment for the specific network settings, also known as operating or performance parameters, such as protocol type, data type, packet size, and traffic usage characteristics, firmware type, operating system type, typical network performance characteristics, and typical, average, peak, and minimum traffic load on the network. For wireless network components, additional factors concerning propagation, signal strength, interference, and noise must be considered.

The preferred embodiment of the invention allows data communication networks to be accurately characterized for performance prediction in a number of novel ways.

First, performance prediction may be based on field measurements from an actual network, where prediction models are formed from some fit to measured data (an empirically-based model). These field measurements may be made manually, or autonomously, using data collectors, or agents, that continually measure and update the specific network performance metrics that are observed within the physical environment. These data collectors are able to measure, or are assigned, specific 3-D position locations within the physical environment, such position locations corresponding to known positions in the computer model which is used to model the physical environment of the

US 6,973,622 B1

15

network, and which are known or which are transmitted to a measurement server. The data collectors may be individuals who manually or automatically record or collect observed network performance such as one or more of the aforementioned performance parameters, or the measurement agents may be software or hardware or firmware applications that run on top of network applications for the purpose of routinely measuring for one of more of the numerous network performance parameters listed previously. The agents may be fixed, or may be portable, and may have position location devices, such as GPS or inertial navigation, or an internal map which is activated by a user, so that the position location of the measurement is sent to a server processor. The agents are presumed to have two-way communication with a server processor that may be collocated or remotely located. Measurements from one or more data collectors are routinely or periodically collected and then transmitted, either by wireless or wired means, or by real-time or stored means, to a server processor which is either collocated, or remotely located, from one or more of the measurement agents. For example, the measurements may be recorded by autonomous agents and then transmitted over a fixed network to a processor that integrates all measurements and computes statistics for observation. The measurement sources have known positions in 3-D, or may not be known and used to form a gross estimate of observed network performance. The collected measurements may be sent in real time, stored and forwarded, or sent as file transfers via many means, such as via email, over the world wide web, via wireless, wired or optical links, or in a storage device. This "in-situ" measurement data is passed, with the 3-D position location when available, to the server, which catalogues and processes the specific measurement information. Using the measurement information from the data collectors, the server is able to provide a predictive model by using knowledge of the physical 3-D environment, and by fusing the many collected inputs into a simplified model of performance that is related to the 3-D physical representation of the world.

In the preferred embodiment of the invention, the server stores and processes the physical location of all measurement devices (where available) as well as all network components and their electrical, logical and technical configuration, while also considering cost and maintenance issues associated with each network component. Using the preferred embodiment, a data communications network can be designed, deployed, tested, predicted, measured, optimized and maintained by collecting the measured data from one or more agents, and processing them at the server to determine a proper prediction engine that allows future network layout with a desired outcome prior to installation. The server engine is able to display the measured results, in a site-specific manner from each measurement agent (that has site-specific information) so that predictions may be compared to measurements on a visual display of a computer or in a stored means (such as an ASCII file comparing predicted versus measured performance parameters).

It is important to note that each measurement agent may be a server, capable of fusing measurement data with the site-specific 3-D layout of the network components and the physical environment. Therefore, each measurement agent may serve as a centralized processor, as well, so that many different physical locations of a particular network may be measured and predicted for performance. Servers may then be collocated or remotely located from the measurement agents, which collect, display, store and use the measurements to form predictive models. In the case of a remote

16

server that receives measurement data from measurement agents, it is possible to remotely monitor, and then predict, the performance of a network that is physically very far from the particular server processor.

The measurement agents may be further controlled or configured by the server processor, so that the agents may be tuned or instructed to perform different types of measurements, such as different packet transmission rates, observation intervals, averaging intervals, protocol types, or other sensible changes which those skilled in the art would conceive for proper network optimization.

A second method for predicting the performance of network parameters is through the use of analytical or simulation methods. These analytical and simulation methods are well known, and relate the physical and electrical characteristics of the network channel to the physical and electrical characteristics of the various network components. Through simulation or analysis, it is possible to determine approximations or bounds on the typical values that one would expect in an actual network configuration of specific components. The present embodiment of the invention allows a user to enter the results of such calculations, so that they are applied as inputs to the prediction model. Therefore, a user of the invention may simply enter "blind" values, based on known methods, as a first guess approach to forming a prediction model of network performance. These first-guess values may then be iterated by the invention, based on feedback from the site-specific measurements of the actual network.

A measured set of data for a typical operating environment with multiple transmitters in a wireless or wired network, are recorded, stored and displayed by the invention, as taught in the previous description about the measurement agents and server processors. Then, some form of best-fit algorithm (minimum mean square, median filter, etc.) may be applied to the predictive models provided in the equations taught below to provide a table look-up for determining proper performance values (e.g. proper values for constants or functions in the performance parameter equations listed below) for a particular site-specific network design. This table look up method allows measured data to be translated into values that may then be used to drive predicted data for all subsequent predictions conducted within the same site-specific 3-D environment in which measurements were made. Alternatively, best guess performance metric values, or best guesses for the functions or constants in the equations listed below, may be fed into the invention, either manually or automatically through a storage means or via a wireless or wired means from a remote or collocated location, for a specific 3-D modeled network environment, wherein the predicted performance at any space or location with the 3-D environment is based on the first, best guess, predictive models. As explained subsequently, these initial best guess, or "blind" models may be based on simulation, analysis, or some combination thereof. The empirically-based predictive models and the initial best guess predictive models may be used in subsequent environments, different from the environment for which measurements or best guesses were made, and the invention allows a catalogue of models to be used easily by the user for subsequent network prediction or design. Measurements of actual network performance may then be overlaid and displayed and stored simultaneously with the network prediction parameters, for rapid comparison. Furthermore, optimization routines compute the best values for minimum error for new predictive models that match the measured network performance within the environment. Thus, the

US 6,973,622 B1

17

invention allows the user to relate empirically-derived predicted performance parameters or initially guessed network performance parameters within a 3-D site specific configuration of the actual installed or contemplated network, using specific information and physical locations about the network devices and by using the models for wired networks and wireless propagation, multipath, and noise. The model techniques for this invention fuse the many factors that impact network performance into simpler models that support prediction and comparison of measured versus predicted network performance for radio/wireless and wired networks. Thus, performance prediction can be ascertained and compared to measured network performance for use in ongoing network deployment.

Furthermore, by comparing measured network performance metrics to predicted metrics, the invention allows new field measurements to update the previous prediction models in a convenient method, which provides a catalogue of models that is stored and displayed to the user either locally or remotely. Alternatively, using the hierarchy of servers, it is possible to use remotely located servers which compute, transmit, or receive such measurements and predictive models for the remote use, display, measurement and storage of model parameters and results. This is particularly convenient for network administrators who wish to monitor the performance and design of networks that are physically distant from the network of interest.

Measurements of a particular device for desired performance criteria is accomplished either by using the measurement software module available in the preferred invention or by importing a log file from another software or hardware measurement tool. The measurement module within the preferred invention allows the measurement of the performance of any specific portion of a communications network using two or more software programs which are installed and run on either sides of a device or devices. These software programs are called agents. By sending test transmissions between two agents across a specific network connection the preferred invention can measure any particular performance criterion. The results of these measurements are stored for a particular portion of the network.

The preferred embodiment of the invention can also import the logfiles of other measurement programs such as traceroute to measure specific links. This functionality allows site-specific measurements made by external programs to be stored site-specifically. This is accomplished by a two-pass method described in patent 09/221,985, "System for Creating a computer model and measurement database of a wireless communication network" by T. Rappaport and R. Skidmore, filed Dec. 29, 1998. To import a logfile a user simply clicks a point in the model of the environment for each data point to assign a location for each point in the logfile.

In performing network performance measurements, especially for wireless data networks, it is important to know the difference in performance for transmission and reception. This is why the preferred invention can measure the transmission and reception components of the average network statistics. To measure the transmission direction, the size of test packets is varied. By changing the size of the packet sent and the size of the packet returned, the transmission and reception statistics can be separated. This allows a network designer to identify problems in transmission that might otherwise be masked by apparently good reception.

Network performance measurements are not useful if the measurements do not mimic the actual data traffic that a network carries. For this reason, the preferred embodiment

18

of the invention is able to mimic the traffic patterns, network protocols and packet characteristics of actual data. Thus, if web browsing performance is being measured, the invention sends small packets from an access terminal to a web server and returns large packets from that server that are typical of text, image and web script file formats. By measuring the performance of such packets, the invention accumulates accurate network statistics for expected web browsing performance.

The measurements of specific traffic types may also be applied to the use of broadcast or multicast packet performance scenarios. The preferred embodiment of the invention is able to measure performance of multiple transmitters or multiple receivers or both of the same packet information. The performance of this type of transmission are different than point to point measurement because shared resources are used more efficiently in broadcast and multicast scenarios. Thus, the ability of the invention to measure network performance statistics for the overall success of the broadcast or multicast transmission and for each individual transmitter and receiver is quite powerful. This ability allows network designers to better choose which transmitters of multicasts might be redundant or which broadcast transmissions are insufficient to reach all the desired receivers.

In some data communications network, the performance of specific pieces of equipment, such as Ethernet Bridges or even a single cable, is hard to measure because it is transparent to the network layer of a data communications network. For this reason, the ability of the invention to determine the performance of a single device through extrapolation is quite useful. The preferred embodiment of the invention is able to use known performance data for specific pieces of network equipment and extrapolate the contribution of other devices in the network. Measuring and extrapolating enough individual hardware and software links can identify the performance of all network devices. The accuracy and reliability of this procedure heavily depends on an accurate and site-specific model of the data communications network, which the invention possesses.

Extending the extrapolation concept of performance evaluation to the software and hardware components of network equipment demonstrates a further capability of the preferred embodiment of the invention. The invention is able to distinguish in some cases between the performance limits due to software and those due to hardware. For example, in a situation where the transmitter and receiver are the same computer, no hardware is actually involved in the transmission. By measuring network statistics in this situation, one can quantify the performance of just the computer software. By comparing the situation where the transmitter and send are the same to a situation where the transmitter and receiver are different computers the performance of just the computer hardware can be identified. Since the performance of the software in either case will be quite similar, the performance of just the hardware in a connection between two computers can be extrapolated by assuming the software will perform similarly in either case.

Extrapolating the performance of individual network components from measured performance metrics can be time consuming. For this reason, the preferred embodiment of the invention is able to read in data results from a plethora of measurement tools, system utilities and network logfiles to a single internal format. The invention is capable of reading in the output of command line utilities such as ping or tcp, the logfiles generated by routers and switches such as tcpdump, or even the logfiles of other commercial measurement programs, and these measurement results are

US 6,973,622 B1

19

stored for use in the predictive engine. The combination of these imported files to a single internal format allows the invention to combine many different measurements and activity logs into a single set of network statistics. This process means the invention requires fewer active measurement campaigns and more diverse and accurate data for better and more accurate network performance modeling.

Accurate, reliable representations of a data communication network require a large number of measured data points. Hence, the preferred embodiment of the invention collects a large amount of data quickly and easily using various methods as described above. The invention does this by providing remote data collection agents, which can be installed on data access terminals or embedded in hardware, software, or firmware within an actual device in the network. The remote data collection agents respond to a server program (the processing server) that controls the measurements made by the remote agent. That is, the remote agent can be directed to make a measurement to or from any other remote agent or processing server using any desired protocol, traffic type, network setting, or configuration. This process does not require any input from a human user at the remote agent's physical location. The agents simply records the data when asked with the correct settings and reports the results back to a server which stores data from all remote agents and other measurement tools. The server can generate a variety of detailed reports and use the data to make predictions about expected network performance in future. Servers can also function as agents. In this manner, servers can be organized in a hierarchy or a distributed fashion. This allows servers to report measurements to one another and make measurements using other agents or servers. A network designer at a server can then use all collected and reported data to identify problem areas such as fairness or poor distribution of broadcast data, or problem times, such as increased network activity at lunch time with a data communications network.

In order to improve the value of measurement data collected, the preferred embodiment of the invention identifies the exact (if possible) or approximate location of a remote agent. As discussed earlier, remote agents in this case can either be controlled by a user at that physical location, or controlled remotely by a server. In the preferred embodiment of the invention, the agent uses information about the network layout to identify an approximate location. Determining the nearest piece of network equipment and associating the approximate location with the precisely known location of that network equipment accomplishes this. This approximate location can be further refined using dead reckoning, clicking on a location in a map, or using the global positioning system, laser range finders or some other positioning device known now or in the future.

The preferred embodiment of the invention is not only capable of accounting for the effects of different hardware, firmware, software and configuration settings, but it can also predict the effects of just the hardware and firmware, just the software, or of a single configuration setting. The ability of the invention to measure and thus adjust empirically-derived predictions for these effects allows the optimization of the data communications network. By predicting the effects of changing any detailed aspect of the data communications network, a user can immediately visualize the effect of a new component or a setting change. This ability allows a user skilled in the art to design an optimal data communications network by continually making changes and observing the prediction changes.

20

We now focus on the details for predicting values for network performance parameters based on knowledge of the 3-D site-specific environment as well as the specific components used in the network design.

The throughput and bandwidth of a network are calculated by the invention as functions of any or all of the following operational parameters which impact performance: distance between transmitter and receiver, physical environment specification, packet sizes, error and source coding schemes, packet overhead, modulation techniques, environment, interference, signal strength, number of users, and for wireless networks, the antenna pattern and type, multipath delay, number of multipath components, angle of arrival of multipath components, radio frequency bandwidth, protocol, coding scheme, and 3-D location. In order to predict the bandwidth and throughput of a network connection, the appropriate functions and constants may be calculated from the listed parameters and then predicted for each location and time desired.

For a wired network, throughput (T) or bandwidth (BW) may be derived from a vendor's specification sheet of a product or device, or may be measured in a special laboratory setting. Alternatively, T or BW may be calculated through analysis or simulation, or may be measured in the field using a number of known devices. These means may be used to determine the proper value for T or BW in a network prediction enging such as contemplated here. A formula for predicting the throughput and bandwidth for a wireless data communications channel is shown in equation 1.

$$T \text{ or } BW = C_1[Ad + Bd^2 + C] + \quad (1)$$

$$C_2[D(RSSI) + E(RSSI)^2 + F] + C_3 \sum_{i=1}^M (G_i P_i + K_i)$$

where T is throughput, BW is bandwidth, d is the distance between a transmitter and a receiver. RSSI is the received signal strength intensity, which is the power level of the signal at the receiver, either in absolute values or in logarithmic values. A, B, C, C₁, C₂, C₃, D, E, F, K_i, are constants or may represent linear or nonlinear functions of one or more physical or electrical parameters, such as physical environment type, packet size, modulation, modem type, or other parameters that relate the physical, electrical, or logical environment of the network. These constants or functions take on specific functional values depending upon if T or BW is being solved for. The value M may denote a particular number of multipath components from a particular transmitter, as determined by propagation analysis of the channel, or the term may denote a combination of important multipath components from a collection of transmitters, where the term "important" is based on antenna pattern, physical environment distances, and other wireless propagation factors which are well known to one skilled in the art and which are explained below. The values of G_i and P_i represent gains and power levels, respectively, for each of M different signal components, which may represent individual multipath components or gross signal components from one or more radiating sources, and K_i represents a finite number of constants or functions for each value of i. Note that G_i, P_i, and the individual K_i may be in logarithmic (e.g. dB) or absolute values. These constants or functions in the above equation may be dependent on distance (d) between transmitter and receiver where d may be the straight-line or actual reflected/diffracted distance of the main signal path between

US 6,973,622 B1

21

the serving transmitter and receiver, 3-D environment, time of observation or observation interval, noise power, packet sizes, coding scheme, number of users, modulation type, interference, and for wireless networks, may include path loss, multipath delay, number of multipath components, angular spread, strength and angle of arrival of received signals, modulation bandwidth, and other physical, electrical and logical settings of particular equipment in the network, and the constants or functions may be calculated analytically, predicted for an initial guess, or solved using best fit methods between measured and predicted performance of actual networks in a site specific environment.

It is important to note that multipath delay, and its effect on network performance prediction and design, may be considered in many ways, as contemplated by this invention and as shown in Equation (1). First, multipath may be considered individually, whereby each multipath component is considered to arrive from each transmitting device, and the methods for modeling multipath are well known and explained in the prior art, and in numerous research works by Rappaport, et. al. from Virginia Tech. Alternatively, gross multipath effects may be modeled as having a worst-case delay (e.g. propagation distance, d) being approximated by the maximum, average, or median length of the specific building or 3-D environment in which the communication network is modeled. Alternatively, spatial considerations may be used by contemplating the antenna patterns of each transmitter or receiver, so that multipath which arrives only in the main beam of each wireless device is considered in the calculation of delay and in network performance in (1). Alternatively, only the strongest one or two or some finite number of transmitters may be considered for multipath propagation delays, whereby only a finite set of transmitters, such as those most closest to the receiver of interest, or those of a certain standard, frequency, or power setting, are considered to radiate multipath energy and produce RSSI values, and from that finite number of transmitters, only the strongest multipath, or the average, maximum, median, or largest few multipath components are considered in computation of delay. Alternatively, if only a finite number of transmitters are considered, methods described above, such as consideration of the physical environment to determine a gross multipath delay from each transmitter, or the use of a particular antenna pattern to determine most important multipath components, may be used to drive the model of multipath and its impact on network performance. Similar approaches may be used to model the received signal strength, RSSI in equation 1.

Note that the constants or functions of equation (1) may be assigned blindly for initial predictions, and then a specific network within the site-specific environment may be measured empirically so that a best-fit (using a minimum mean square error approach or some other well known method) may be used to assign values for the constants or functions in (1). Note that in (1), the distance (d) may be based on true physical distance from the 3-D site specific model of the environment, or may actually represent a relative distance ratio, where the physical distance between two points is referenced to a convenient close-in free space reference distance, as is customary for propagation predictions, and is taught in (Rappaport, "Wireless Communications, Principle & Practice, Prentice-Hall, 1996).

Propagation delay for network data is predicted for wired networks, where components are interconnected by wire (either fiber or metal wire) by dividing the distance traveled by the propagation speed of the electrical, electromagnetic or optical signals in the device, which are used to transmit

22

the data. For instance, data in a fiber optic cable travels at a speed 2×10^8 meters per second due to dielectric properties of the cable, which affect the photons in a fiber optic cable that are used to transmit the data. Such photons move at the speed of light in glass, which is less than the free space propagation speed. Thus, if the cable is 200 meters long the transmission delay is equal to 1×10^{-6} seconds. By using the site-specific method of modeling the complete network within the present invention, it is possible for the user to simultaneously visualize the network as configured in the environment and see a display of delay and predicted or measured performance of delay within the cable within the 3-D environment. Additionally, using a tool tip mouse cursor or some other pointing means, or using a pull down menu, or by simply viewing the display device which the invention is implemented on, various network performance metrics, as well as stored data from the Bill of Materials and parameters of interest may be visualized or stored.

Predicting the propagation delay for a wireless portion of a data communications network is more difficult than wired networks due to the fact that multiple transmitter sources, such as access points in a Bluetooth network, IEEE 802.11b, or wireless ATM network may be transmitting simultaneously. Furthermore, as mentioned previously, multipath interference can create echoes that may or may not be equalized depending on the specific network equipment used at the wireless receiver or transmitter. However, the same calculation model used for wired networks may be used, with the additional consideration of multipath delay terms, and propagation losses or gains, due to specific multipath components, as shown in Equation (1). This additional consideration of multipath delay is needed to account for the fact that wireless data does not always travel in a straight line, and that physical objects can diffract, reflect, absorb, and scatter radio energy. Thus, to calculate the transmission delay of a wireless link in a data communications network, the distance between the transmitter and the receiver is divided by the propagation speed (3×10^8 meters per second) of a wireless communications link and then added to the multipath delay introduced by the indirect paths taken from transmitter to receiver as is shown in equation 2.

$$T_p = \frac{d}{3 \times 10^8 \text{ m/s}} + \tau_d \quad (2)$$

Where T_p is the propagation delay in seconds, d is the distance between the transmitter and the receiver in meters, and τ_d is the multipath delay in seconds. Predicting the multipath delay is performed using well-known raytracing techniques or based on angle of arrival, or signal strength values, or by making estimated based on the physical model of the 3-D environment.

Transmission delay is directly calculated from the bandwidth of a connection using the number of bits transmitted. To calculate transmission delay, the number of transmitted bits is divided by the bandwidth. This calculation is identical for wired and wireless channels but must be performed separately for each network device. The formula is illustrated in equation 3.

US 6,973,622 B1

23

$$T_t = \frac{\# \text{ of bits}}{BW} \quad (3)$$

Where T_t is the transmission delay time in seconds, # of bits are the number of bits in the transmission or packet and BW is the bandwidth of the network link in bits per seconds.

Processing delay must be calculated for each device separately within a network. Processing delay is the time required for a network device to process, store, and forward the data bits that are applied to a network device. Alternatively, processing delay may be the time required for a source to produce a meaningful data stream once it is instructed to do so. Processing delay is known to be zero for devices that do not perform any processing, such as passive network components like cables, antennas, or splitters. Processing time may depend on the packet size, protocol type, operating system, vendor, firmware, hardware, and software versions or configurations, and the type of device and the current computing load on the device. To predict the processing delay of any device it is necessary use a model that accounts for all of these effects. These models may be measured in the field, measured in a test facility, obtained from vendors, or derived from analysis or simulation.

Queuing delay is only applicable to devices that transmit data from multiple users or multiple connections. The queuing delay of a device is the amount of time a particular packet must wait for other traffic to be transmitted. It is difficult to predict the queuing delay of a particular connection because it depends on the amount of traffic handled by a particular device. For this reason, queuing delays can be predicted using a statistical random variable based on the expected performance of the device and/or the expected traffic load. Alternatively, average, median, best or worst case, or some other linear or nonlinear weighting of queuing delay times as defined by the device specifications, or as measured, simulated, or computed by analysis, may be used to calculate a predicted queuing delay time.

Packet latency, round-trip times and handoff delay times are all based on propagation, transmission, processing, and queuing delay times. To accurately predict packet latency and round trip time, the propagation, transmission, processing and queuing delay times must be summed for all network devices in a particular network link and adjusted using the particular traffic type, packet size, and protocol type. For instance, packet latency is the time required for a packet to travel from transmitter to receiver. To predict packet latency for a particular link the propagation, transmission, processing and queuing delay times must be calculated using the specific network connection, traffic type, and packet size for the one-way transmission of a packet.

Round trip times are calculated similarly, except for the transmission and reception of a packet and the return of the acknowledging packet. Thus, to predict the round trip time, the invention takes into account the original packet size and the size of the acknowledging packet as well as the effects of the specific network connection, protocol and traffic type on the propagation, transmission, processing and queuing delays.

Handoff delay times are based on the propagation, transmission, processing and queuing delays involved in two separate wireless access points coordinating the change of control of a wireless device from one access point to another. These delays result because the two access points must transmit data back and forth to successfully perform a

24

handoff. Thus, the prediction of handoff delay time is similar to the prediction of the packet latency time between the two access points. To predict the handoff delay time, the invention calculates the propagation, transmission, processing and queuing delays for the link between the two access points. The invention then adjusts for the specific number of transmissions required and the size of the data, which must be sent to successfully perform a handoff.

When predicting bit error rates, the invention considers wired and wireless error rates. Wireless networks operate in much more hostile electrical environments than their wired counterparts and their interconnections are significantly more difficult to model and, until this invention, practical networks have not successfully been modeled using specific, accurate physical and electrical models of multiple transmitters, multiple interferers, noise sources, and network components within a 3-D site-specific environment. This invention uses 3-D site specific representations of the environment for specific network implementations that are able to consider both wired and wireless networks, and considers physical locations, electrical specifications and attributes of all radiating sources and their antenna systems in a real-world 3-D environmental model. Wireless networks are prone to data errors much more so than wired channels, due to the impact of multipath propagation, multiple transmitters, and noise, as described previously. The fact that radio propagation and noise is more random than for fixed wired networks must be considered for practical design, and is modeled in this invention. For wired channels, bit error rates are simply a measure of the electrical, optical and electromagnetic parameters of a connection and are predicted using a statistical random variable, such as a Gaussian or Poisson random distribution, or other sensible distribution or algorithm known now or in the future, and this random variable is overlaid about the average, median, or typical performance of the network component or network subsystem. The network device or subsystem may include a single wireless node, such as a router or switch, or a complete interconnection of various routers, hubs, switches, wireless access points, and wireless client/server devices that communicate with the network. The network may be wired, wireless, or a combination thereof.

Many performance metrics of a device or a network subsystem, such as Frame Error Rate, Bit Error Rate, or Packet Error Rate, as well as other performance parameters such as throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time may be either derived from a specification of the equipment, may be calculated analytically within the invention or inputted into the invention, or may be measured a priori in advance to using the invention. That is, specific parameters of operation, known as operating parameters or equipment parameters, such as those listed previously, can be either measured or predicted through equipment specifications provided by vendors. Alternatively, they may be measured in-situ by a user or research facility, for proper modeling and input into the invention. Alternatively, they may be calculated based on some known analytical model that contemplates interconnection of devices so that a performance model and operating parameters may be computed. The statistical random variable to model network performance within the invention can be dependant on the electrical, optical and electromagnetic characteristics of each device such as voltage levels, power levels, impedance,

US 6,973,622 B1

25

and operating frequencies, or can be generated using a typical observed (measured) value for each network device. For instance, copper wire can be modeled as having a bit error rate of 1 error in 10^6 or 10^7 bits transmitted. Once measured and characterized a single initial time, a single component or a string of components within a network may be modeled repeatedly by the invention, so that network performance models.

Wireless performance parameters, however, are dependant on many more factors than wired bit error rates. For this reason, the invention predicts wireless bit error rates based on the environment, distance between transmitter and receiver, number and types of partitions obstructing the transmission, time, 3-D position, packet size, protocol type, modulation, radio frequency, radio frequency bandwidth, encoding method, error correction coding technique, multipath signal strengths and angle of arrival, and multipath delay. As a result, the calculation of the predicted bit error rate is performed using constants or functions to convert from previously measured or known channel and network equipment performance metrics to an expected bit error rate. A formulation for predicting the bit error rate, frame error rate or packet error rate directly for a data communications channel is shown in equation 4, and is identical to equation 1.

$$BER, PER, \text{ or } FER = C_1[Ad + Bd^2 + C] + \quad (4)$$

$$C_2[D(RSSI) + E(RSSI)^2 + F] + C_3 \sum_{i=1}^M (G_i P_i + K_i)$$

where BER is bit error rate, FER is the frame error rate, PER is the packet error rate, d is the distance between a transmitter and a receiver. RSSI is the received signal strength intensity, which is the power level of the signal at the receiver. A, B, C, C_1 , C_2 , C_3 , D, E, F, K_i , are constants or linear or non linear functions with different values depending on which of BER, FER, and PER is being calculated. The value M may denote particular number of multipath components from a particular transmitter, or may denote a combination of important multipath components from a collection of transmitters, where the term "important" is based on antenna pattern, physical environment distances, and other wireless propagation factors which are well known to one skilled in the art and which are explained within this disclosure. The each of M values of G_i and P_i represent gains and power levels, respectively, of different signal components, which may represent individual multipath components or gross signal components from one or more radiating sources, and may be in logarithmic or linear values of power. The variables G_i and P_i and each one of the M number of K_i values may be in logarithmic (e.g. dB) or absolute values. These constants in the above equation are dependant on distance (d) between transmitter and receiver where d may be the straight-line or actual reflected/diffracted distance of the main signal path between the serving transmitter and receiver. As explained in the text surrounding equation (1), distance may be straight-line distance, or may be modeled from the gross characteristics of the environment, such as the maximum, average, or median length of the 3-D environment. As with equation (1), equation (4) may consider the distance d as the actual physical distance, or as a relative distance referenced to a close-in reference.

Frame error rates, packet error rates and packet drop rates can all be calculated from bit error rates or predicted directly

26

using the same method as for a bit error rate as described above or as modeled in equation 4. To perform these calculations the invention uses information stored in the site-specific Bill of Materials about the packet size, frame size and the protocol in use, and uses a site-specific propagation and interference modeling technique, such as that utilized in the SitePlanner product by Wireless Valley Communications, Inc.

In wireless networks, modeling the combined effects of all the various sources of errors is extremely difficult. Not only does modulation and specific error and source coding techniques impact the wireless network performance, but so does the impact of antennas, multipath, noise, voice over IP or wireless ATM concatenation methods, modem design of particular wireless modem makers, and the specific RF distribution system used to connect wired and wireless devices. The ability to model such varied effects can be done by allowing field measurement of specific in-situ network performance as explained earlier. By conducting a walk-through or a drive test whereby a mobile receiver is operated and network performance parameters are measured within the site-specific environment, it is then possible to determine best fits for particular modem manufacturers, applying concepts described in equation 1.

Bandwidth delay products can be calculated by the invention directly using information about any or all of the environment, three dimensional position, protocol type, multipath delay, packet sizes, radio frequency, radio frequency bandwidth, coding, number, strength and angle of arrival of multipath components, signal strength, transmission, propagation, processing and queuing delay, bit error rate, packet error rate, and frame error rates. Alternatively the invention can calculate the bandwidth delay product indirectly using previously predicted values. A bandwidth delay product is calculated by multiplying the bandwidth of a certain network device by the total delay introduced by that device. Thus, the formula is illustrated here in equation 5:

$$BWD = \frac{BW}{T_{net}} \quad (5)$$

Where BWD is the bandwidth delay product, BW is the bandwidth and T_{net} is the total delay introduced.

The invention uses statistical models of the consistency of data communications network hardware to predict packet jitter and quality of service (QoS). Both of these performance criteria are measures of the reliability of a network to provide consistent data arrival times. Thus, to calculate the QoS or jitter of a connection, the invention uses formulas which include any or all of the environment, three dimensional position, protocol type, multipath delay, packet sizes, radio frequency, radio frequency bandwidth, coding, number, strength and angle of arrival of multipath components, signal strength, transmission, propagation, processing and queuing delay, bit error rate, packet error rate, frame error rate, throughput, bandwidth, and bandwidth delay product. The formulas include constants or functions, which relate the above variables in general to the variation in the arrival time of data and in specific to the QoS and packet jitter of a connection. The present embodiment of the current invention uses equations (1) or (4) to determine QoS and packet jitter for a data communications network.

The preferred embodiment of the invention predictions consider the effects of not just the site specific, floor plan, building layout, terrain characteristics and RF characteris-

US 6,973,622 B1

27

tics, but also the effects of the particular network hardware, firmware and software in the network. The invention allows the network to be modeled down to the settings and locations of the individual data communications devices, using the Bill of Materials discussed earlier. The prediction of network performance statistics takes these settings into account. This means that different transport level protocols (such as TCP or UDP), different protocol settings (such as packet and buffer sizes), the data bandwidth (in bits per second), physical layer transmission methods including modulation techniques (such as QPSK or FHSS), coding schemes (such as CCK or trellis codes), transport media (such as copper, fiber optic cable or wireless connections) and specific frequency bands are taken into account by the invention. These aspects are in addition to the consideration of the location and wireless specific criteria, which includes transmitter-receiver distance (T-R distance), the propagation environment, interference, path loss, number of users sharing the RF resources, multipath delay, the number of multipath components and their strengths and angle of arrival, the ratio of coherent to incoherent power, and the RF bandwidth (in Hz). All of these variables may produce results which may be mapped into the form of equation (1) or (4).

The predictions of the preferred form of the invention consider the characteristics of the data communications network users. Information such as the type of data communications traffic the users offer to the network, the number of users, and the usage patterns over time, are stored in a location specific manner in the invention. That is, points can be placed which represent individual users and the traffic offered by that user or areas in which the characteristics of a group or pool of users can be assigned. The invention takes these points and areas of user traffic into account when making predictions of network performance criterions. This means that if large numbers of users are found in an area covered by access points that are able to adapt to heavy usage, the invention is able to accurately predict the performance of these (or any other) conditions. This is only possible because of the accurate, location specific model of the data communication network. Additionally, since the preferred form of the invention tracks usage patterns of users over time, the resulting measurements may be used by a server processor to form table look-up values for the constants or functions of Equations (1) or (4). Different values of constants or functions for Equations (1) or (4) may be found to predict the performance of the network at different times of day. This is an important aspect of a data communication network prediction model because real networks have peak usage times and lulls in which usage is lower. By tracking the usage of a data communications network over time, the preferred form of the invention can determine if the network will have difficulties at certain times.

In a communications network, the capacity is always a scaled version of the theoretical maximum possible capacity, and the impact of various users, and their propagation characteristics, message sizes, as well as the network characteristics, all combine to bound or limit the capacity that an individual user sees on a network. Consider a network that has, as a bottleneck, a particular component or device which has a maximum rating of T_{max} bits per second. This component bounds the maximum possible throughput of the network. Consider that capacity represents the capacity or throughput of a device or network (defined as T or Capacity), where T(x,y,z,t)=T_{max}[γ], where γ is a scaling factor that fuses many different, complicated physical, electrical, and logical conditions into a simple value that ranges between 0 and 1. When gamma is 0, there is no capacity.

28

When gamma is 1, there is maximum capacity. Note that T is a function of 3-D positioning in the network, as well as a function of time. For a particular user, the goal of a network predictive model is to predict the capacity, as a function of 3-D position and as a function of time. Thus, T[x,y,z,t] will range between 0 and T_{max}.

The load put on to a data communications network impacts the capacity of an individual user. The number of users and the usage patterns of each user affect the capacity of each user in a data communications network. The preferred embodiment of the invention allows a network designer to see the effects of network loading on the important network statistics, by measuring the instantaneous traffic conditions with the measurement agents as described above. It is possible to determine in-situ capacity measurements through other means, such as observation from network equipment or reporting mechanisms built into hardware or software products. By forming a table look-up of the specific capacity results, as a function of 3-D site-specific location, as well as the time of day, the invention builds a measurement-based predictive model for capacity. These measurements may be used to form a model of capacity, as now presented.

The invention contemplates the fact that the scaling factor on capacity (or throughput), is a function of the instantaneous number of users of the network, the maximum number of simultaneous users of the network, the average and maximum packet size used by users of the network, and for many other factors that are modem or network or vendor or protocol specific. Also, in the case of a wireless network, the multipath propagation effects, the propagation distances between the user and the wireless access points, and the received signal levels are factors that limit capacity. In addition, constants or functions that fuse the impact of modulation, equalizations, impulse noise, and other factors, are used in the invention.

Thus, capacity or throughput of a network is modeled by

$$\text{Capacity} = C_1 [Ad + Bd^2 + C] + \quad (6)$$

$$C_2 [D(RSSI) + E(RSSI)^2 + F] + C_3 \sum_{i=1}^M (G_i P_i + K_i)$$

where the constants or functions of (6) take on similar properties as described for equations (1) and (4). Furthermore, the entire equation (6) may be scaled by K/U_{max}, where K is the instantaneous number of users on the network, and U_{max} denotes the maximum number of simultaneous users possible.

Handoff delay times are potential problems in wireless data communication networks. A handoff occurs in wireless data networks when a user moves out of range of one access point and into range of another access point. In this situation, the first access point must pass the responsibility of delivering data to the wireless user to the second access point. If the two access points are too far apart, there will not be enough time for a wireless data network user to be handed off from one access point to another and file transfers can fail. The invention predicts where handoffs will occur and the possibility of handoff failures due to incompatible network settings at two different access points by using site-specific time dependent measurements, and fitting them into a form of equation (1), (4) or (6). Then, a table look up method is used to determine prediction models for handoff times as a function of spatial positioning and time of day.

US 6,973,622 B1

29

The concept of optimization is a key aspect of the invention. The preferred invention is highly effective at allowing one skilled in the art to quickly improve the performance of an existing data communications network by comparing measured performance parameters with predicted values that are derived and stored in the invention. The process of using measurements to improve predictions is called optimization and is illustrated in FIG. 6, FIG. 7, and FIG. 8. The method for optimizing a network using just measurements is shown in FIG. 6, just predictions in FIG. 7, and a combination of measurements and predictions in FIG. 8. The process of optimizing a data communications network is accomplished by comparing, through numerical, visual, or some other means, the predictions and measurements of performance criteria such as throughput, bandwidth, quality of service, bit error rate, packet error rate, frame error rate, dropped packet rate, packet latency, round trip time, propagation delay, transmission delay, processing delay, queuing delay, network capacity, packet jitter, bandwidth delay product and handoff delay time for various site-specific locations and particular times of day. By changing the hardware used in the network, or changing the locations of hardware or the configuration of that hardware, firmware, or software which controls each device within the network, one skilled in the art can improve the performance of the network. These performance improvements can be implemented and viewed by repeating predictions of the performance criteria after site-specific equipment changes to the network have been made in the 3-D model of the network. Continuing this process allows one skilled in the art to optimize the performance of a network to achieve an efficient data communications network.

Using this information, the preferred embodiment of the invention can make recommendations for the areas of the network to upgrade or reconfigure. The invention can also use SNMP protocol communications or other protocols to actually implement these changes. That is, a network designer could identify problems in a data communications network through prediction, whereby the prediction of performance criteria of the data communications network is calculated using known measurement data and the configuration and expected performance of all data communications hardware in the data communications network. The predicted performance criterion is stored and displayed visually and numerically in a location specific, three-dimensional model of the environment. Then, the designer can use the invention to identify a solution to the problems that are apparent by viewing the prediction results, either by following the inventions recommendations for changes or making the designers own change. After simulating the predicted outcome, the network designer can then direct the invention to update all the relevant settings of the equipment with the changes the designer has just used in a prediction. The designer could then use the tool to measure the results of these changes using the measurement features of the invention.

While this invention has been described in terms of its preferred embodiments, those skilled in the art will recognize that the invention can be practiced with considerable variation within the scope of the appended claims.

What is claimed is:

1. A method for analyzing and adjusting a wireless communications network, comprising the steps of:
generating or using, with a computer or server, a computerized model of a wireless communications network within a physical space in which said wireless communications network is deployed, said computerized

30

model providing a site specific representation of one or more of a floor plan, building layout, terrain characteristics, or RF characteristics, said computerized model identifying locations within said physical space of one or more components used in said wireless communications network, said computerized model having modeled attributes for at least one each of said one or more components;

receiving, at said computer or server, measurement data from one or more measurement collectors or agents located in said physical space, said one or more measurement collectors or agents being the same or different from one or more of said one or more components used in said wireless communications network;

predicting, at said computer or server, one or more performance metrics for said wireless communications network, wherein predictions are made based on said modeled attributes for said at least one of said one or more components, and said measurement data from said one or more measurement collectors or agents; and changing settings or configurations of at least one component of said wireless communications network based on instructions sent from said computer or server.

2. The method of claim 1 wherein said site specific representation is three dimensional.

3. The method of claim 1 wherein said data collection measurement collectors or agents are portable or fixed.

4. The method of claim 1 further comprising the step of affixing said measurement collectors or agents permanently within said physical space.

5. The method of claim 1 wherein said performance metric predicted in said predicting step is selected from the group consisting of throughput, error rates, packet latency, packet jitter, symbol jitter, quality of service, security, coverage area, bandwidth, bit error rate, packet error rate, frame error rate, dropped packet rate, queuing delay, round trip time, capacity, signal level, interference level, bandwidth delay product, handoff delay time, signal-to-interface ratio, signal-to-noise ratio, physical equipment price, and cost information.

6. The method of claim 1 wherein said measurement data received in said receiving step obtained manually.

7. The method of claim 1 wherein said measurement data received in said receiving step obtained autonomously.

8. The method of claim 1 further comprising the step of storing said measurement data.

9. The method of claim 1 further comprising the step of updating said computerized model.

10. The method of claim 9 wherein said step of updating includes the steps of:

specifying components from a plurality of different modeled components which are to be used in said communications network, said modeled components including descriptions and attributes of a specific component; and specifying locations within said physical space for a plurality of different components in said computerized model.

11. The method of claim 10 wherein said step of updating further includes the step of specifying an orientation for at least one component specified in said first specifying step at said location specified in said second specifying step.

12. The method of claim 1 wherein said computerized model identifies orientations of said components at said locations within said physical space and said predicting step utilizes said orientations.

13. The method of claim 1 wherein said computerized model includes one or more objects which create noise or

US 6,973,622 B1

31

interference, said noise or interference being an attribute of said one or more objects which are factored in said predicting step.

14. The method of claim 1 wherein said one or more performance metrics predicted in said predicting step are predicted in a forward direction in said wireless communication network.

15. The method of claim 1 wherein said one or more performance metrics predicted in said predicting step are predicted in a reverse direction in said wireless communication network.

16. The method of claim 1 further comprising the step of specifying data transfer protocol, and wherein said predicting step uses a specified data transfer protocol as a factor in predicting said one or more performance metrics.

17. The method of claim 1 further comprising the step of specifying a network loading for said wireless communications network, and wherein said predicting step uses a specified network loading in predicting said one or more performance metrics.

18. The method of claim 1 further comprising the step of storing or visualizing data representing comparisons of measurements with predictions.

19. The method of claim 1 further comprising the step of storing or visualizing data representing either or both logical connections of network components or physical locations of network components.

20. A system or apparatus for analyzing and adjusting a wireless communications network, comprising:

a computer or server for generating or using a computerized model of a wireless communications network positioned within a physical space, said computerized model providing a site specific representation of one or more of a floor plan, building layout, terrain characteristics, or RF characteristics, said computerized model identifying locations within said physical space of one or more components used in said wireless communications network, said computerized model having modeled attributes for at least one of said one or more components;

one or more measurement collectors or agents operating or operational within said physical space which send measurement data to said computer or server, said computer or server predicting one or more performance metrics for said wireless communications network based on said measurement data and said modeled attributes for said at least one of said one or more components, and said computer or server can send instructions to one or more components of said wireless communications network which cause settings or configurations of at least one component to be changed.

21. The system or apparatus of claim 20 wherein said site specific representation is three dimensional.

22. The system or apparatus of claim 20 wherein said measurement collectors or agents are portable or fixed.

23. The system or apparatus of claim 20 wherein said measurement collectors or agents are permanently affixed at within said physical space.

24. The system or apparatus of claim 20 wherein said performance metric predicted by said computer or server is selected from the group consisting of throughput, error rates, packet latency, packet jitter, symbol jitter, quality of service, security, coverage area, bandwidth, bit error rate, packet error rate, frame error rate, dropped packet rate, queuing delay, round trip time, capacity, signal level, interference

32

level, bandwidth delay product, handoff delay time, signal-to-interface ratio, signal-to-noise ratio, physical equipment price, cost information.

25. The system or apparatus of claim 20 further comprising a storage device for storing said measurement data.

26. The system or apparatus of claim 20 wherein said computerized model is stored on at least one server, wherein said at least one server is the same or different from said computer or server.

27. The system or apparatus of claim 26 wherein said computerized model is stored on a plurality of servers, and said plurality of servers can communicate with each other.

28. The system or apparatus of claim 27 wherein said plurality of servers have a heirarchical relationship to one another.

29. The system or apparatus of claim 26 further comprising at least one portable client device, said at least one portable client device can communicate with said at least one server.

30. The system or apparatus of claim 28 wherein said system includes a plurality of portable client devices.

31. The system or apparatus of claim 20 further comprising a storage medium or display for, respectively, storing or visualizing data representing comparisons of measurements with predictions.

32. The system or apparatus of claim 20 further comprising a storage medium or display for, respectively, storing or visualizing either or both logical connections of network components or physical locations of network components.

33. A method for analyzing and adjusting a wireless communications network, comprising the steps of:

generating or using, with a computer or server, a computerized model of a wireless communications network within a physical space in which said communications network is deployed, said computerized model providing a site specific representation of one or more of a floor plan, building model, terrain characteristics, or RF characteristics, said computerized model identifying locations within said physical space of one or more components used in said wireless communications network, said computerized model having modeled attributes for at least one of said one or more components;

downloading or inputting files of measurement data to said computer or server, where said measurement data is obtained from said physical space or from said wireless communications network;

predicting or providing a one or more performance metrics for said wireless communications network based on said measurement data and said modeled attributes for said at least one of said one or more components; and

changing settings or configurations of at least one component of said wireless communications network based on instructions sent from said computer or server.

34. The method of claim 33 wherein said measurement data is obtained from measurement collectors or agents that are either portable or fixed.

35. The method of claim 33 further comprising the step of storing or visualizing data representing comparisons of measurements with predictions.

36. The method of claim 33 further comprising the step of storing or visualizing data representing either or both logical connections of network components or physical locations of network components.

37. A site specific method for analyzing and adjusting a communications network, comprising the steps of:

US 6,973,622 B1

33

generating or using, with a computer or server, a computerized model of a communications network positioned within a physical space, said computerized model providing a site specific representation of one or more of a floor plan, building layout, terrain characteristics or RF characteristics, said computerized model identifying locations within said physical space of one or more components used in said communications network, said computerized model having modeled attributes for at least one of said one or more components

receiving, at said computer or server, measurement data from one or more measurement collectors or agents located in said physical space, said one or more measurement collectors or agents being the same or different from one or more of said one or more components used in said communications network;

predicting, using said computer or server, one or more performance metrics for said communications network, wherein predictions are made based on said measurement data and said modeled attributes for at least one of said one or more components;

changing settings or configurations of at least one component of said communications network based on instructions sent from said computer or server.

38. The method of claim 37 wherein said site specific representation is three dimensional.

39. The method of claim 37 wherein said measurement collectors or agents portable or fixed.

40. The method of claim 37 further comprising the step of affixing said measurement collectors or agents permanently within said physical space.

41. The method of claim 37 wherein said one or more performance metrics predicted in said predicting step are selected from the group consisting of one or more performance metrics are selected from radio signal strength intensity, connectivity, network throughput, bit error rate, frame error rate, signal-to-interference ratio, signal-to-noise ratio, frame resolution per second, traffic, capacity, signal strength, throughput, error rates, packet latency, packet jitter, symbol jitter, quality of service, security, coverage area, bandwidth, server identification parameters, transmitter identification parameters, best server locations, transmitter location parameters, billing information, network performance parameters, C/I, C/N, body loss, height above floor, height above ground, noise figure, secure coverage locations, propagation loss factors, angle of arrival, multipath components, multipath parameters, antenna gains, noise level reflectivity, surface roughness, path loss models, attenuation factors, throughput performance metrics, packet error rate, round trip time, dropped packet rate, queuing delay, signal level, interference level, quality of service, bandwidth delay product, handoff delay time, signal loss, data loss, number of users serviced, user density, locations of adequate coverage, handoff locations or zones, locations of adequate throughput, E_c/I_o , system performance parameters, equipment price, maintenance and cost information, user class or subclass, user type, position location, all in either absolute or relative terms.

42. The method of claim 37 wherein said measurement data received in said receiving step is obtained manually.

43. The method of claim 37 wherein said measurement data received in said receiving step is obtained autonomously.

44. The method of claim 37 further comprising the step of storing said measurement data.

34

45. The method of claim 37 further comprising the step of updating said computerized model.

46. The method of claim 45 wherein said step of updating includes the steps of:

specifying components from a plurality of different modeled components which are to be used in said communications network, said modeled components including descriptions and attributes of a specific component; and specifying locations within said space for a plurality of different components in said computerized model.

47. The method of claim 46 wherein said step of updating further includes the step of specifying an orientation for at least one component specified in said specifying components step at said location specified in said specifying locations step.

48. The method of claim 37 wherein said computerized model identifies orientations of one or more of said one or more components at said locations within said physical space and said predicting step utilizes said orientations.

49. The method of claim 37 wherein said computerized model includes one or more objects which create noise or interference, said noise or interference being an attribute of said one or more objects which are factored in said predicting step.

50. The method of claim 37 wherein said one or more performance metrics predicted in said predicting step are predicted in a forward direction in said communication network.

51. The method of claim 37 wherein said one or more performance metrics predicted in said predicting step are predicted in a reverse direction in said communication network.

52. The method of claim 37 further comprising the step of specifying data transfer protocol, and wherein said predicting step uses a specified data transfer protocol as a factor in predicting said performance metric.

53. The method of claim 37 further comprising the step of specifying a network loading for said communications network, and wherein said predicting step uses a specified network loading in predicting said one or more performance metrics.

54. The method of claim 37 further comprising the step of storing or visualizing data representing comparisons of measurements with predictions.

55. The method of claim 37 further comprising the step of storing or visualizing data representing either or both logical connections of network components or physical locations of network components.

56. A site specific system or apparatus for analyzing and adjusting a communications network, comprising:

a computer or server for generating or using a computerized model of a communications network positioned within a physical space, said computerized model providing a site specific representation of one or more of a floor plan, building layout, terrain characteristics or RF characteristics, said computerized model identifying locations within said physical space of one or more components used in said communications network, said computerized model having modeled attributes for at least one of said one or more components;

one or more measurement collectors or agents positioned within said physical space which obtain and send measurement data to said computer or server, said computer or server predicting one or more performance metrics for said communications network based on said measurement data and said modeled attributes for said at least one of said one or more components, and said

US 6,973,622 B1

35

computer or server can send instructions to one or more components of said communications network which cause settings or configurations of at least one component to be changed.

57. The system or apparatus of claim 56 wherein said site specific representation is three dimensional.

58. The system or apparatus of claim 56 wherein said measurement collectors or agents are portable or fixed.

59. The system or apparatus of claim 56 wherein said measurement collectors or agents are permanently affixed at locations within said physical space.

60. The system or apparatus of claim 56 wherein said one or more performance metrics selected from the group consisting of one or more performance metrics are selected from radio signal strength intensity, connectivity, network throughput, bit error rate, frame error rate, signal-to-interference ratio, signal-to-noise ratio, frame resolution per second, traffic, capacity, signal strength, throughput, error rates, packet latency, packet jitter, symbol jitter, quality of service, security, coverage area, bandwidth, server identification parameters, transmitter identification parameters, best server locations, transmitter location parameters, billing information, network performance parameters, C/I, C/N, body loss, height above floor, height above ground, noise figure, secure coverage locations, propagation loss factors, angle of arrival, multipath components, multipath parameters, antenna gains, noise level reflectivity, surface roughness, path loss models, attenuation factors, throughput performance metrics, packet error rate, round trip time, dropped packet rate, queuing delay, signal level, interference level, quality of service, bandwidth delay product, handoff delay time, signal loss, data loss, number of users serviced, user

36

density, locations of adequate coverage, handoff locations or zones, locations of adequate throughput, E_c/I_0 , system performance parameters, equipment price, maintenance and cost information, user class or subclass, user type, position location, all in either absolute or relative terms.

61. The system or apparatus of claim 56 further comprising a storage device for storing said measurement data.

62. The system or apparatus of claim 56 wherein said computerized model is stored on at least one server which may be the same or different from said computer or server.

63. The system or apparatus of claim 62 wherein said computerized model is stored on a plurality of servers, wherein said plurality of servers can communicate with each other.

64. The system or apparatus of claim 63 wherein said plurality of servers have a heirarchical relationship to one another.

65. The system or apparatus of claim 62 further comprising at least one portable client device that can communicate with said at least one server.

66. The system or apparatus of claim 64 wherein said system includes a plurality of portable client devices.

67. The system or apparatus of claim 56 further comprising a storage medium or display for, respectively, storing or visualizing data representing comparisons of measurements with predictions.

68. The system or apparatus of claim 56 further comprising a storage medium or display for, respectively, storing or visualizing either or both logical connections of network components or physical locations of network components.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 6,973,622 B1
APPLICATION NO. : 09/668145
DATED : December 6, 2005
INVENTOR(S) : Rappaport et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 30, line 7 change "each of" to --of--

Column 30, line 38 change "interface" to --interference --

Column 30, line 42 insert "is" after "step"

Column 30, line 44 insert "is" after "step"

Column 32, line 2 change "interface" to --interference--

Column 32, line 21 insert "or apparatus" after "system"

Column 32, line 48 delete "a"

Column 33, lines 36, 37 delete "one or more performance metrics are selected from"

Column 35, line 14 delete "one or more performance metrics are selected from"

Signed and Sealed this

Fifth day of September, 2006

A handwritten signature in black ink, appearing to read "Jon W. Dudas". The signature is stylized with a large, looped initial "J" and a cursive "Dudas".

JON W. DUDAS
Director of the United States Patent and Trademark Office